

TOSCA

A Comprehensive Brain-based Cognitive Architecture

September 27, 2006

This document is an initial draft of the design of the TOSCA architecture. Our goal above all else was to be comprehensive – to provide a design for a complete mind that is grounded in the brain. Of necessity in a project this ambitious, our design is still incomplete in many details, but where it is incomplete we have included the known constraints that a final design will have to meet. Completing the architecture will require an exploration, which we will pursue using the TOSCA framework (see Appendix). The TOSCA framework is a software infrastructure that directly supports expressing the connectivity constraints of the brain, while providing the flexibility to rapidly develop alternative implementations of the functional modules and state variables that make up the TOSCA architecture.

TOSCA Participants (funded under BICA Phase I)

BICA Thrust A:

- Michigan (John Laird, Richard Lewis, Thad Polk, Doug Pearson (Three Penny))
- MIT (Cynthia Breazeal, Linda Smith (Indiana), Larry Barsalou (Emory))

BICA Thrust B:

- Dartmouth (Richard Granger, Carey Priebe (Johns Hopkins), Anna Tsao (Algotek))
- Harvard (Stephen Kosslyn, Bruce Draper (CSU))
- Rutgers (Mark Gluck)

1. Introduction and overview

Our design of TOSCA starts at the brain system and circuits levels. In developing an initial version of TOSCA, we've chosen to intentionally abstract away from much of the complexity of the brain. Many brain systems include multiple subsystems that are extremely complex in their own right (e.g., vision and hearing within sensory systems) and the sophisticated computational mechanisms underlying these systems are important, but we first need to define the “forest” – the overall architecture with the major pieces and how they fit together, before we get to the “trees”. This is purely a tactical decision to get us started and we fully plan to greatly expand the systems and subsystems in TOSCA in the future. Our strategy is to include those neural systems that we consider most important in constructing an initial functional architecture that provides end to end behavior.

The document has the following structure. In Section 2 we identify the innovative claims of our design that distinguish it from other approaches to systems and circuit level models of the brain. Section 3 describes the high-level structure of the architecture in terms of the basic (repeated) architectural loops connecting major cortical and subcortical regions. This basic loop includes perception, categorization (with clustering), access to memories, internal and external actions and action selection, as well as feedback. Section 4 then lays out in more detail the major theoretical commitments concerning the operation of each of the major brain subsystems and their inter-connections, and initial assumptions about representation, time-course, and algorithms. In Section 5, we identify the key emergent functional properties that derive from the integration of the components; many of these properties concern the multi-faceted nature of learning in the system.

2. Innovative claims

Some of the basic assumptions underlying TOSCA are shared by other system and circuit level models of the brain. Nevertheless, we list them here because they are quite different from the assumptions underlying most standard symbolic AI systems and many cognitive models:

- Brain as control system, but no central controller.
- Asynchronous, parallel, distributed processing.
- Multiple internal memories at different temporal and spatial scales.
- Continuous learning throughout subsystems.
- Specialization/localization of function and content, but unifying cross-function/cross-content processing principles.

TOSCA itself is *distinguished* by an innovative set of claims concerning the computational structures and algorithms that give rise to cognition in the brain:

- Perception:
 - Top-down processing plays a dominant role in perception.
- Categorization and Identification:
 - Internal representations are learned by building clusters of sequences of clusters across all sensory modalities.
- Internal knowledge representation, learning, memory, and use:
 - Knowledge is represented as distributed multi-modal structures, but with local structure.
- Mental operation selection:
 - Mental (internal) operations use the same basic brain structures as external action selection.
- Action learning:
 - Intrinsic reward based reinforcement learning drives learning of external and internal action selection.
- Integration:
 - By integrating these different learning mechanisms in a complete architecture, the resulting whole should be more powerful than the sum of the parts. Incorporating intrinsic rewards allows useful learning to occur from simple exploration without any explicit task. Learning over mental operations as well as motor actions leads to the development of cognitive skills as well as motor skills. And combining reinforcement learning with clustering over state and action representations makes possible the acquisition of complex skills contingent upon very abstract features.

The TOSCA architecture is more than a sum of its parts, and it has novel characteristics and principles of interaction that further distinguish it from other BICAs. These characteristic and principles emerge from the details of TOSCA and thus will be presented in Section 5 after the complete architecture has been described.

3. Overall structure: The basic loop

Our design starts with the human brain, which consists of evolutionarily recent forebrain circuit designs (telencephalic circuits) layered on top of preserved ancient (e.g., reptilian) circuits, with the new designs accounting for more than 90% of the volume of the human brain. There are four primary divisions of telencephalic forebrain (cortex, striatal complex, hippocampal formation, amygdala nuclei), and many subdivisions (e.g., anterior vs posterior cortex, distinct cortical layers, local circuits, striatal components, hippocampal fields CA1, CA3, dentate gyrus, subiculum, ...), each with its own cell types and local circuit design layouts, thus presumably each conferring unique computational properties.

There is (perhaps surprisingly) a single large-scale architecture that organizes all telencephalic components. For almost any given region of posterior cortex, there is a corresponding region of anterior cortex (e.g., the frontal eye fields, connected to posterior visual cortical areas), as well as corresponding regions of striatum, pallidum and thalamus, connected in register. These complementary cortical and subcortical regions are connected in a characteristic pattern: reciprocal connections between posterior and anterior cortex, converging anterior and posterior cortical projections to a related region of striatum, which in turn connects (via pallidum and thalamus) back to the same region of anterior cortex. This overall “systems circuit” is by far the largest coherent loop in the mammalian brain, and it is repeated for multiple regions of posterior cortex, with dedicated regions corresponding to individual sensory modalities, as well as non-cortical telencephalic regions including components of hippocampus and amygdala, connected with dedicated regions of striatum and anterior cortex.

We are faced with a difficult problem in describing the design of TOSCA. The operation of a specific component is important, but the interaction among components is equally (or even more) important. Moreover, we have interactions between groups and loops of components. Our approach is to initially focus on the basic loop of behavior from perception to action and describe the primary neural systems that participate in that loop. This will leave out some structures that play a less central role in the basic loop. In going through the loop we often give a cursory description of a component because our goal is to build up the big picture, emphasizing interactions.

At the highest level, the initial version of the TOSCA architecture will attempt to tightly integrate the most important neural systems found in the brain. These neural systems are described in detail in Section 4, and are labeled below with their corresponding subsection. The descriptions in parentheses summarize the main computational functions we attribute to each system.

- 4.1. Sensory systems (Low-level vision and audition)
- 4.2. Specific thalamocortical (core) circuits (Clustering)
- 4.3. Non-specific thalamocortical (matrix) circuits (Sequencing)
- 4.4. Cortico-cortical circuits (Bottom-up associations, top-down control, and self-organizing maps)
- 4.5. Cortico-hippocampal circuits (Episodic memory)
- 4.6. Cortico-striatal circuits (Action selection)
- 4.7. Dopamine reward circuits (Reinforcement learning)
- 4.8. Cortico-amygdala circuit (Emotion, State-dependent storage & retrieval)

Figure 1 presents some of the major pathways among the structures involved in these neural systems. Many others that play a less central role in the basic operation of TOSCA are omitted. In contrast to many traditional cognitive architecture where functionality maps directly to architectural components, in TOSCA (and the brain), it is the circuits and loops through multiple neural structures from which functionality emerges.

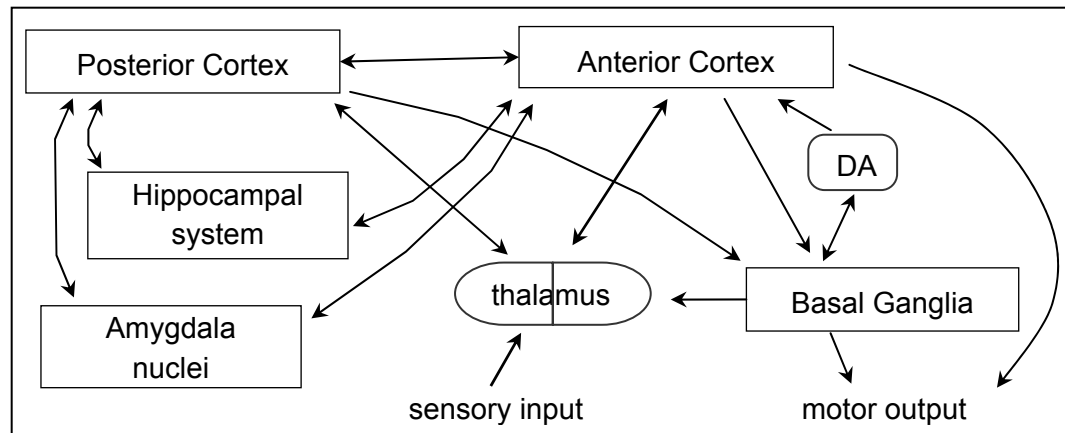


Figure 1: Major pathways within and between neural systems underlying TOSCA.

This pattern of connections is consistent with the following processing loop:

1. Posterior cortex potentially receives any of five types of input: sensory input (via thalamus), episodic memory (via hippocampal system), top-down control signals (from anterior cortex), emotion (amygdala), and bottom-up associations (from other parts of posterior cortex).
2. Specific (core) thalamocortical loops cluster on patterns of activity, recognizing familiar input patterns.
3. Non-specific (matrix) thalamocortical loops encode (& retrieve) sequences of clusters, producing a representation of what is expected to come next.
4. Cortico-cortical connections within posterior cortex modify and elaborate the internal state, generating a more complete hierarchy of clusters of sequences of clusters that is organized topographically within cortical areas.
5. Cortico-cortical projections to anterior cortex propose specific intentions, which could be motor actions or mental actions (e.g., setting cues for episodic memory retrieval, setting goals in working memory, maintaining or attending to specific information).
6. Corticostriatal loops select among competing actions based on the (learned) values associated with each action in the current context/state. Multiple actions can be selected in parallel, based on the parallel structure of frontostriatal loops.
7. State-action values are modified by the midbrain dopamine system in a way that realizes reinforcement learning algorithms (strengthening state-action associations that lead to long-term reward).
8. Motor actions are passed on to the motor output systems for execution whereas mental actions provide top-down control signals to posterior cortex.

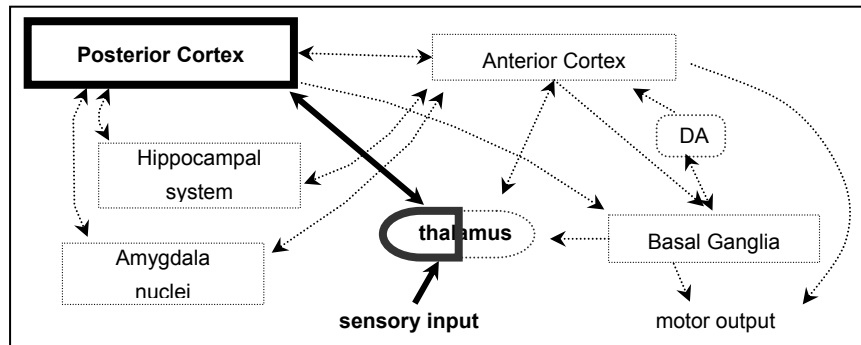
4. Major brain subsystems

This section describes the major brain subsystems that will be modeled by TOSCA. These are not isolated regions of the brain but are instead circuits involving multiple brain regions. For each of these subsystems we first present the underlying anatomical structures in the brain, with an accompanying figure that highlights those structures in Figure 1 that contribute to the circuit. This is followed by a description of the physiological operation of the structure. We then discuss the derived computational functionality – this is the core of what we will implement in software. This is followed by a description of how this subsystem interacts with the rest of the system.

4.1 Sensory systems (low-level vision and audition)

Anatomical structure

The early visual system contains among the most intensively studied regions of the brain, including the retina, dorsal lateral geniculate nucleus of thalamus, and primary visual cortex (V1), as well as lateral occipital cortex



(LOC). The early auditory system, also heavily studied, consists of the cochlea, auditory brainstem nuclei, medial geniculate nucleus of thalamus, and primary auditory cortex (A1).

In both cases, the feed-back lines from cortex back to thalamus outnumber the feed-forward lines from thalamus to cortex by an order of magnitude, strongly suggesting the powerful role played by top-down cortical modulation of early structures, as described below.

Physiological operation

Although many issues remain to be understood in these systems, much agreement has been reached in the scientific literature about the nature of initial processing throughout these structures. Many independent laboratories have arrived at shared assumptions about the processing of these regions, which extract initial motion, edge, depth, and primitive shape information in visual areas, and edge contours, formants, and detailed frequency information in auditory areas. Importantly, as mentioned, these areas receive an order of magnitude more input back from cortex than they provide output to cortex. As might then be expected, these areas respond powerfully to extensive feedback signaling from cortex, enabling a range of attentional and gating operations. Rather than a view in which signals pass from the periphery into cortex, the system is instead a highly active perceiver, controlling and predicting inputs throughout the perceptual process and directing the paths by which inputs arrive. The extensive published literature on these mechanisms has been oft-replicated and documented, enabling past work to serve as a solid starting point for further downstream processing in TOSCA.

An illustrative tale can be told in this regard. Initial work on cochlear front-end mechanisms yielded extensive data on the processing being carried out by cochlea and other early

structures. At one time there was an extensive program of study aimed at carefully replicating cochlear function in silicon devices, with the possible goal of treatments for the deaf. During this work, experiments were performed on deaf subjects with simple, low-resolution frequency filters, bypassing the cochlea and plugged directly into the auditory nerve. To the considerable surprise of many scientists in the field, these experiments had the following result: initially, little intelligible sound got through to the subjects – but after a “practice” period averaging 2-3 weeks, the subjects “learned” to interpret the sounds; learned so well that many could understand spoken language over a telephone, a notoriously low-bandwidth and noisy conduit. It emerged that the downstream areas (e.g., auditory cortical regions) learned to interpret the inputs as though they were conveyed by normal biological apparatus, even though the inputs were measurably and substantially inferior to the signals that would have been sent by a real cochlea. These experimental findings strongly support the lesson that top-down cortical control can overcome even extensive loss of fidelity in front-end structures.

Derived computational functionality

As a result of the widespread agreement in the literature on early visual processing, work in TOSCA can take advantage of these shared findings and rely on widely agreed-upon standard “off the shelf” mathematics and software systems corresponding to front end processing (Gabor filters, edge detectors, motion detection). It is anticipated that little new work will be done on these feature extraction systems.

In particular, one method in testing for Tosca is a difference of Gaussians (DoG) filter applied to multi-scale information channels (“pyramids”) to produce a salience function, followed by attention windows centered on the local maxima of the salience function. In our case we model covert as well as overt attention, and thus can select multiple attention windows per frame, and we select the scales as well as locations of attention windows, as suggested by data showing that covert attention may create both translational and scale invariance. We model feature extraction as a process of voting in parameter spaces that reflect color and non-accidental edge properties, thus corresponding to high-dimensional feature vectors describing each attention window. Most importantly, Tosca includes top-down salience channels that indicate how well an image matches scaled versions of predicted objects, as described in subsequent sections.

As per the previous discussion, widespread consensus is accruing that most of the “heavy lifting” in both auditory and visual understanding is being carried out downstream, and via “top-down” modulatory control by cortical structures of early sensory areas. The TOSCA team brings extensive longstanding background expertise on the primacy of top-down processing in vision.

Systems

Early vision and auditory systems are strongly interconnected with downstream perceptual object, category, and location recognition structures, as described. It is also noteworthy that most visual and auditory areas project to multimodal areas, in which visual and auditory and even somatosensory information become intermingled. Thus these downstream higher perceptual and cognitive areas apparently must share functions common to different modalities. This principle is exploited throughout the TOSCA architecture.

4.2 Specific thalamocortical (core) loops (clustering)

Anatomical structure

Projections from cells in thalamic “core” nuclei synapse on target neurons in all cortical layers to some extent but predominantly in deep layer III and (in granular cortical regions) in layer IV, as well as on the

apical dendrites of layer VI neurons. These afferents, which preserve topographic organization, are often described as the primary input to sensory neocortical regions, though quantitative neuroanatomical studies report that these thalamic inputs constitute a very small percentage of the total set of afferents to cortical layer IV cells: for instance, inputs from dorsal lateral geniculate nucleus (dLGN) neurons comprise less than 6% of the synaptic contacts onto layer IV target cells in primary visual cortex. Projections from a given thalamic core region extend to a cortical area roughly 0.5 – 1.0 mm wide, somewhat larger than the size of physiologically delineated functional columns. Layer VI axons project back topographically to the thalamic core cells from which they receive inputs, as well as to the overlying portion of the nucleus reticularis (NRt) covering the target core cells. NRt in turn generates GABAergic projections to these thalamic core cells. The result can be depicted by highlighting a subset of the connections that occur in thalamocortical circuits as in Figure 2.





Figure 2: Thalamocortical Circuits

Physiological operation

Here we describe the simplified steps that occur in response to normal inputs. Peripheral inputs activate thalamic core cells which in turn participate in topographic activation of middle cortical layers, e.g., ear → cochlea → auditory brainstem nuclei → ventral subdivision of medial geniculate (MGv), or corresponding thalamic core nucleus (“core” in the figure), → primary auditory cortex (A1), layer IV → layer II-III → layer VI → N.Ret → MGv (core).

In the event that a fixed input is being focused on, i.e., that a stream of inputs is not arriving, then this loop will recur. (In the next section, it will be seen that in response to a stream of inputs, other thalamocortical mechanisms become engaged, interrupting operation of the core loop and employing elements of both core and matrix loops).

As the core loop recurs in response to fixed input stimuli, a series of physiological responses occurs. The superficial (layer II-III) cells that are most responsive to a given input will in turn activate neighboring inhibitory (red) cells, which then in turn inhibit all the excitatory cells in the region. Thus the response to an input is a relatively brief discharge from only the most responsive excitatory cells followed by silence induced by lateral inhibition. As the synaptic contacts onto the responding cells become strengthened via LTP, those cells become increasingly probable responders even to slightly different spatial input patterns. Thus those superficial cells that initially respond to a particular input pattern become increasingly responsive not only to that input but also to a range of similar inputs (those that share many active lines; i.e., small Hamming distances from each other), such that similar but distinguishable inputs will come to elicit identical patterns of layer II-III cell output, even though those inputs would have given rise to slightly different output patterns in the absence of LTP.

These learning-based (LTP-based) effects can be simply characterized in terms of the formal statistical operation of clustering, in which sufficiently similar inputs are placed into a single category or cluster. This is further discussed in the next section, on functional implications.

Immediately following this response from superficial layer neurons, those cells activate deep layers (V and VI; see figure). Output from layer VI initiates feedback activation of nucleus reticularis (N.Ret) (Liu and Jones 1999) which in turn inhibits the core thalamic nucleus. Since, as described above, topography is preserved through this sequence of projections, the portions of the core nucleus that become inhibited will correspond topographically to those portions of L.II-III that were active. On the next cycle of thalamocortical activity, the input (assumed as above to be a relatively fixed unchanging input) will arrive at the core nucleus against a background of inhibitory feedback from N.Ret, which has been shown to last for hundreds of milliseconds (Cox et al., 1997; Zhang et al., 1997). Thus the predominant component of the next input to cortex is only the uninhibited remainder of the input, whereupon the same operations as before are performed. Thus the second cortical response will consist of a quite distinct set of neurons from the initial response, since many of the input components giving rise to that initial response are now inhibited. This process of inhibition and distinct selected responses continues until the feedback inhibition at N.Ret diminishes (roughly 500 – 1000 msec).

Derived computational functionality

Analysis of the sequence of responses in computational models has shown clustering and successive sub-clustering of inputs. The first cycle of response identifies the input's membership in a general category of similar objects (e.g., flowers); the next response (a fraction of a second later) identifies its membership in a particular subcluster (e.g., thin or fat flowers); then sub-sub-clusters, etc. Thus the system repetitively samples across time, differentially activating specific target neurons at successive time points, to discriminate among inputs (see, e.g., Kilborn 1996; Rodriguez et al., 2004).

Table 1: Simplified Thalamocortical Core Algorithm

```

for input X
  for C ∈ win(X, W)
     $W_j \leftarrow W_j + k(X - C)$ 
  end_for
 $X \leftarrow X - \text{mean}(\text{win}(X, W))$ 
end_for
where
  X = input activity pattern (vector); W = layer I synaptic weight matrix;
  C = responding superficial layer cells (col vector); k = learning rate parameter;
  win(X, W) = column vector in W most responsive to X before lateral inhibition [ $\forall j, \max(X \cdot W_j)$  ]

```

The method can be characterized as an algorithm (Table 1). Analysis reveals the algorithm's time and space costs. The three time costs for processing of a given input X are: i) summation of inputs on dendrites; ii) computation of “winning” (responding) cells C ; iii) synaptic weight modification. For n learned inputs of dimensionality N , in a serial processor, summation is performed in $O(nN)$ time, computation of winners takes $O(n)$ time, and weight modification is $O(N \log n)$. With appropriate parallel hardware, these three times reduce to $O(\log N)$, $O(\log n)$, and constant time respectively, i.e., better than linear time. Space costs are similarly calculated: given a weight matrix W , to achieve complete separability of n cues, the bottom of the constructed hierarchy will contain at least n units, as the leaves of a tree with $\log Bn$ hierarchical layers, where B is the average branching factor at each level. Thus the complete hierarchy will contain $\sim n[B/(B-1)]$ units, i.e., requiring linear space to learn n cues (Rodriguez et al., 2004).

These costs compare favorably with those in the (extensive) literature on such methods (Rodriguez et al., 2004). Elaboration of the algorithm has given rise to families of computational signal processing methods whose performance on complex signal classification tasks has consistently equaled or outperformed those of comparable methods (Coultrip and Granger, 1994; Kowtha et al., 1994; Granger et al., 1997; Benvenuto et al., 2002; Rodriguez et al., 2004).

Systems

The thalamocortical core loop is part of the overall thalamocortical loop, which includes the matrix circuit, discussed in the following section. Taken together, thalamocortical loops are the primary circuit in the brain, engaged in every cortical region, and in turn participating in cortico-striatal and cortico-limbic (hippocampal and amygdala) circuits, each described later.

4.3 Non-specific thalamocortical (matrix) loops (sequencing)

Anatomical structure

Projections from cells in thalamic “matrix” nuclei predominantly connect in layer I, chiefly on the apical dendrites of neurons from layers II, III, and V. These projections have been referred to as “nonspecific,” in i.e., broad

and diffuse in contrast to the more topographic projections from “core” nuclei in thalamus (Lorente de No 1938; Killackey and Ebner 1972, 1973; Herkenham 1986; Jones 1998). It has consistently been confirmed that matrix cells projecting to a given cortical area receive projections back from layer V of that cortical area without intervening NRt contacts (Conley and Diamond 1990; Rouiller et al 1991; Bourassa and Deschenes 1995; 1998). This portion of the thalamocortical circuit can be illustrated by highlighting a subset of the connections that occur in thalamocortical circuits as in Figure 3.


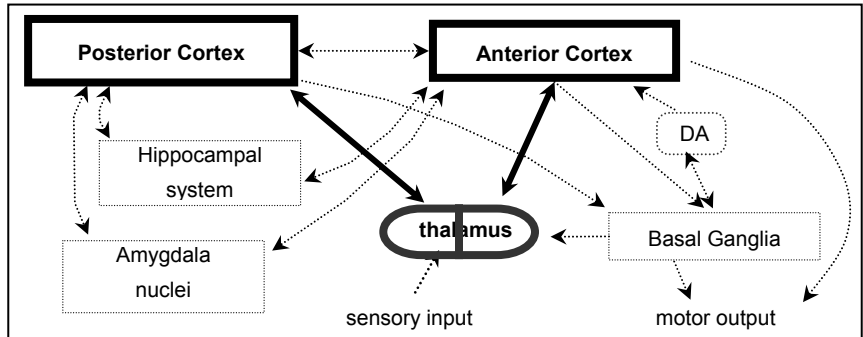


Figure 3: Thalamocortical Matrix Circuit

Physiological operation

Unlike the core loop, the matrix circuit receives no inputs from peripheral signals. Only after cortex is activated by inputs via the core loop, is the matrix loop activated. In particular, once the superficial layer cells in cortex respond to an input, their output activates not just layer VI as described in 4.2 but also layer V, which sends diffuse (non-topographic) feedback to matrix nucleus Mt, which in turn projects back up to layer I of cortex. Their non-topographic nature means that these projections do not retain any neighbor relations that may obtain among inputs.

(This loop, like that of the core circuit, is timed via endogenous “clocks”: synchronous activity of wide regions of cortex (modulated in part by ascending systems affecting the periodic responsivity of inhibitory cells) makes the probability of excitatory cell spiking lower during peak inhibition and higher during inhibitory troughs. Moreover, the average

time course of excitatory and inhibitory potentials (15 msec and 50 msec, respectively), and the time constants of dendrites, severely limit the temporal precision with which a target neuron can “read” differences among slightly different spike trains (Magee 2000).)

Cortical pyramidal cells preferentially respond to onsets and offsets, i.e., transitions among inputs in all cortical areas studied (somatosensory: Peterson et al. 1998; auditory: Recanzone et al. 2000; visual: Rols et al 2001; Bair et al 2002).

The activation of layer V in rapid sequence via activation by superficial layers (in response to each element of a sequence) and via activation by Mt (corresponding to feedback from previous element in the sequence) selects responding cells sparsely from the most activated cells in the layer (Coultrip et al., 1992) and selects synapses on those cells sparsely as a function of the sequential pattern of arriving inputs. Thus synapses potentiated at a given step in layer V correspond to the input occurring at that time step together with orthogonalized feedback arising from input just prior to that time step (Aleksandrovsky et al. 1996; Rodriguez et al. 2004).

Derived computational functionality

The same steps as those described in section 4.2 obtain, but in response to time-varying inputs, a different effect is produced: that of “chaining” the elements in the input sequence via the “links” created due to layer V activity from coincident inputs corresponding to current and prior input elements. As in the operating rule described by Granger et al. (1994), the sparse synaptic potentiation enables the cells in layer V to act as a novelty detector, selectively responding to those strings that have previously been presented. Whereas superficial layer cells in the model respond to any of a number of sufficiently similar inputs (the “clustering” effect described earlier), the deep layer cells respond only to the input sequences that have actually occurred previously, due to the orthogonalizing input from Mt combining with superficial layer input. Thus the layer V activation patterns even for very similar input sequences will be very different from each other, or, put differently the probability of two similar input sequences eliciting similar sequences of layer V patterns is low.

Table 2: Simplified Thalamocortical Matrix Algorithm

```

for input sequence X(L)
    for C ∈ TopographicSuperficialResponse(X(L))
        for V(s) ∈ C ∩ NNtResponse(X(L-1))
            Potentiate( V(s) )
        NNt(L) ← NontopographicDeepResponse(V)
    end_for
end_for
where L = length of input sequence;
C = columnar modules activated at step X(L);
V(s) = synaptic vector of responding layer V cell,
NNt(L)= response of nonspecific thalamic nucleus to feedback from layer V.
```

As before, the method can be characterized as an algorithm (Table 2). Rodriguez et al. (2004) showed that the space costs grow linearly with the number of sequences stored, for an assumed fixed acceptable rate of collision errors.

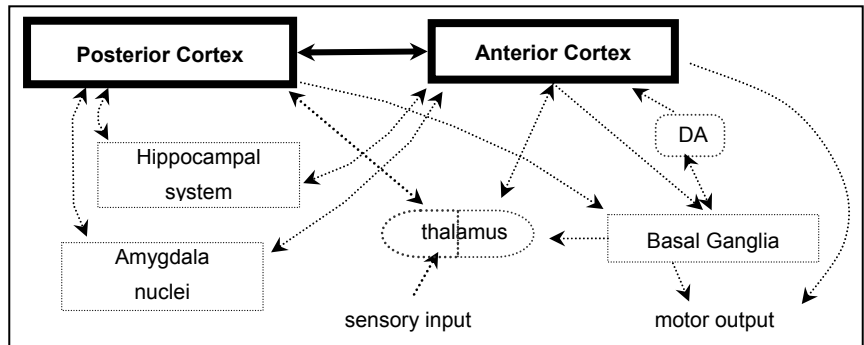
Systems

The thalamocortical loops are part of the overall cortico-cortical and cortical-subcortical systems-level organization of the telencephalic model. The primary representations, hierarchically nested sequences of categories, are elaborated in various ways via these interactions. As will be seen in Section 5, these representations underlie content all the way from perception to language.

4.4 Cortico-cortical circuits (bottom-up associations, top-down control, self-organizing maps)

Anatomical structure

The proposed architecture of cortico-cortical circuits in TOSCA is based on a few general features of cortico-cortical circuits: (1) Distinct cortical areas based on cytoarchitectural differences, (2) massive, bidirectional connectivity within and between cortical areas including substantial bidirectional projections between frontal cortex and posterior cortex, and (3) topographic projections: nearby cells tend to project to nearby targets. TOSCA also incorporates assumptions about the micro-circuitry within cortical columns (interacting with thalamic nuclei) which play a critical role in the emergence of clustering and sequencing (see sections 4.2 and 4.3).



The proposed architecture of cortico-cortical circuits in TOSCA is based on a few general features of cortico-cortical circuits: (1) Distinct cortical areas based on cytoarchitectural differences, (2) massive, bidirectional connectivity within and between cortical areas including substantial bidirectional projections between frontal cortex and posterior cortex, and (3) topographic projections: nearby cells tend to project to nearby targets. TOSCA also incorporates assumptions about the micro-circuitry within cortical columns (interacting with thalamic nuclei) which play a critical role in the emergence of clustering and sequencing (see sections 4.2 and 4.3).

Physiological function

Distinct cortical areas perform quite different functions. At the highest-level, sensory information is processed in posterior neocortex (consisting of occipital, parietal, and temporal neocortex) while anterior cortex (frontal cortex) is primarily involved in processing actions and intentions. At a finer grain-size, there are on the order of 50-100 distinct cortical areas (or more depending on how you divide it up) performing quite different functions. Many of these cortical areas are organized topographically with nearby cells exhibiting similar receptive fields. Cortical representations of multimodal stimuli/concepts involve a large population of active neurons distributed across multiple cortical areas. Processing across cortical areas is strongly interactive as the activity within a cortical area can be strongly influenced both by bottom-up influences (e.g., perceptual processing in sensory cortex) and top-down influences (e.g., attentional influences from anterior areas).

Computational function

At the most coarse level, the initial version of TOSCA will distinguish posterior cortex, which we assume represents the current state, from anterior cortex, which we assume

represents intentions (both motor and mental actions as well as goals). Projections from posterior cortex to anterior cortex will encode associations between specific features of the state and taking (or not taking) specific actions (or more generally, adopting a specific intention which may be an abstract goal that requires many actions to achieve). Conversely, projections from anterior cortex to posterior cortex will encode a mental operation by specifying how a given action should change the state. Simple examples include exciting part of the posterior state representation in order to maintain it (working memory), focusing attention on some particular feature of the state, activating or manipulating a mental image in sensory cortex, among many others. Of course, projections within posterior cortex are also capable of changing state. We assume these projections correspond to better learned, more automatic associations whereas frontal representations correspond to more controlled, deliberate intentions.

As previously discussed, the central representation in TOSCA will be sequences of clusters (of sequences of clusters...). Clusters naturally arise from thalamocortical loops which perform a kind of competitive learning: neurons whose receptive field best matches the current input "win" and their receptive field is modified to be closer to that input. This is a standard approach in connectionist modeling and is known to lead to receptive fields that represent category prototypes.

We adopt the additional assumption that learning affects the receptive fields of neurons that are spatially near the "winning" cells. This assumption is quite plausible under the assumption that nearby cells tend to cooperate (e.g., excite each other) rather than compete. This is the critical assumption underlying all self-organizing map (SOM) models and leads to the kind of topographically organized networks that are ubiquitous in cortex. What this means for TOSCA is that clusters that are represented nearby in cortex will tend to represent similar stimuli/concepts and will tend to project to nearby targets. Such topographic organization naturally supports similarity-based generalization under the assumption that cortical representations correspond to population codes rather than grandmother cells. To see this, consider what happens when an association is learned between one cluster corresponding to a feature of the state and another cluster corresponding to a potential action to take when that feature is present. Each representation corresponds to a large population of nearby cells with similar receptive fields (a population code). Learning the association between them corresponds to strengthening the connections between the two populations. Similar, but slightly different, features of the state will activate an overlapping population of cells as the original feature and, as a result, the new features will be partially associated with the same action. As a result, the architecture will be able to choose actions that are generally appropriate based on states it has never experienced, as long as those states are similar to states it has experienced.

Systems

The cortical system will interact with all the other major subsystems in TOSCA: thalamus (clustering and sequencing), hippocampal system (episodic memory), corticostriatal circuits (action selection), midbrain dopamine system (reinforcement learning), amygdala (reciprocal priming). These interactions are described in the other parts of section 4.

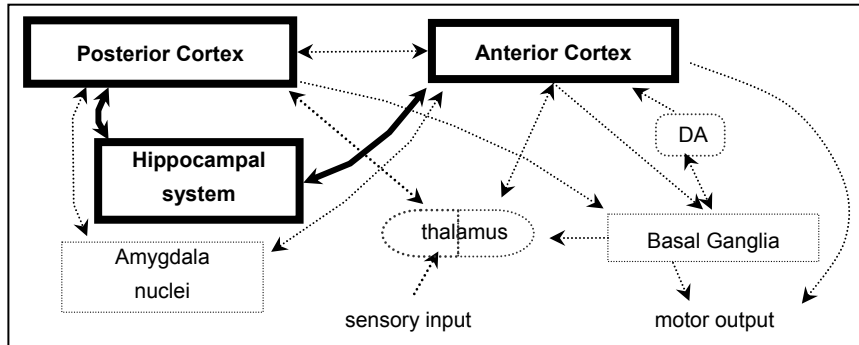
4.5 Cortico-hippocampal circuits (episodic memory, spatiotemporal relations)

Anatomical structure

As described below in more detail, our network model of cortico-hippocampal circuits for learning and memory will include modules corresponding to the

dentate gyrus (DG), CA3 and CA1 fields of the

hippocampus proper, and superficial and deep entorhinal cortex.



Entorhinal cortex. The entorhinal cortex contains six layers that, for simplicity, can be divided into “superficial” (layers I-III) and “deep” (layers V-VI) EC. The superficial layers receive highly-processed multimodal sensory input from neocortex (primarily via perirhinal and postrhinal cortex). Principal neurons in the superficial layers include pyramidal neurons (in layer III) and stellate cells (in layer II). The stellate cells project via the perforant path to DG and CA3, while the pyramidal cells project to CA1 (and subiculum). The superficial layers also contain a large number of GABAergic interneurons that exert a widespread inhibitory control over the output of principal cells. The deep layers receive input from CA1 (and subiculum) and project back to the same neocortical areas that provided input to the superficial layers. There is also a projection from deep to superficial EC that causes both excitation and feedforward inhibition (van Haeften et al., 2003). Pyramidal cells in the deep layers show graded persistent firing (over 5 minutes) which could allow for reverberating circuits (superficial EC to hippocampus to deep EC to superficial EC) to maintain stimulus representations across short delays (Frank & Brown, 2003).

Hippocampal Formation. The hippocampus includes a dentate gyrus (DG) layer, a CA3 layer, and a CA1 layer. Connections from DG to CA3 and from EC to CA1 are topologically organized. Each stellate neuron in EC contacts a subset of the possible postsynaptic targets in DG and in CA3. Each neuron in CA3 contacts a subset of the possible postsynaptic targets in CA3 and CA1.

Physiological function

EC neurons receive external input representing highly pre-processed multimodal sensory information from cortex. They will be modulated by interneurons providing both feedback and feedforward inhibition; for simplicity. Strong inhibitory processes and local circuit feedback in the EC cause representational compression, implementing representational clustering function proposed by Myers et al. (1995). Deep EC neurons form the principal output of the hippocampal region back to cortex and also project to principal cells in superficial EC.

Computational function

In our implementation of TOSCA, we will follow the widely accepted hypothesis that the hippocampal region plays a critical role in the acquisition of new memories, particularly

rapidly-acquired memories for autobiographical events, sometimes collectively called episodic memory (e.g. Squire, 1987; Squire et al., 2004). As a starting point we plan to incorporate our previous neural network modeling of hippocampal region processing in the larger architecture (Gluck & Myers, 1993, 2001; Myers & Gluck, 1994). This model assumes that the hippocampal region develops new stimulus representations that encode contextual and stimulus-stimulus regularities. Specifically, we found that known features of the anatomy and physiology of EC (sparse activation of principal neurons, dense inhibition, and local plasticity mechanisms) give rise to the compression of redundant features in the input. This model accounted for data showing that latent inhibition and sensory preconditioning, which depend on compressing together the representations of CS and context and/or co-occurring cues, survive selective hippocampal lesion but are impaired after EC or broad hippocampal-region damage (Myers et al., 1995). We will adopt this same model in the initial version of TOSCA. We will also follow our previous modeling in assuming that the hippocampal layer forms a compact code for the whole situation in which the organism finds itself (what we call the 'ensemble'; Murnane, Phelps, & Malmberg, 1999). Such representations will form the basis of episodic memory in TOSCA.

Systems

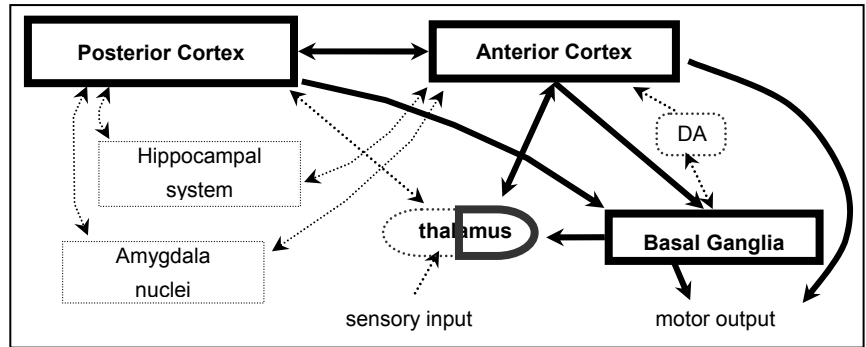
Interactions between the hippocampal system and other neural systems will play a crucial functional role in TOSCA. At the highest level, the hippocampal system will constantly be encoding and storing compressed representations of the current state (as represented in posterior cortex). When similar states are encountered in the future, they will activate the previously stored compressed representation which will in turn reinstantiate information from the previously stored state in posterior cortex. Once this information is represented in posterior cortex, it can influence which actions/intentions are proposed and selected. Furthermore, we envision corticohippocampal loops in TOSCA storing and retrieving temporal sequences of events that have been experienced. Specifically, each event in a sequence could provide cues that lead to retrieval of the next event in the sequence. In this way, the hippocampal system could be used to replay a sequence of events from the past. Doing so could be potentially very valuable to the agent, because it would make it possible to plan ahead and predict likely future events that may improve its present decision making.

The interaction between the hippocampal system and anterior cortex could provide another crucial functionality for TOSCA. Recall that one critical assumption of the architecture is that it learns how and when to perform *mental* operations as well as motor actions. That is, the same learning algorithms will be used to reinforce rewarding actions, whether they are mental actions or motor actions. The initial design of TOSCA will exploit this strategy in order to learn how best to exploit its episodic memory system. For example, TOSCA should be able to learn when the mental act of attempting an episodic memory retrieval is likely to lead to long-term reward. Similarly, it should learn when episodic storage is called for. Indeed, the agent should even be able to learn what retrieval cues to set in posterior cortex in order to retrieve memories that are likely to help in deciding how to act. Put simply, TOSCA should be able to learn how to use its episodic memory most effectively in addition to learning episodic memories themselves.

4.6 Cortico-striatal circuits (action selection)

Anatomical structure

The basal ganglia (BG) are a set of interconnected, sub-cortical nuclei which form a complex network of loops integrating cortical, thalamic and brainstem information (Alexander et al 1986). There are three main pathways from the cortex, through the BG, and back to the cortex (Figure 3).



The striatum is the input nucleus of the direct pathway. It projects directly to the output nuclei of the BG, the Globus pallidus interna (GPe) and Substantia nigra pars reticulata (SNr). The output nuclei project back to the cortex via the thalamus, with the input returning to the same cortical module that provided the excitation to the striatum.

The striatum also has a second pathway to the output nuclei, the indirect pathway. This two step inhibitory pathway provides delayed excitation to the same area of the output nuclei that the striatum inhibited via the direct pathway.

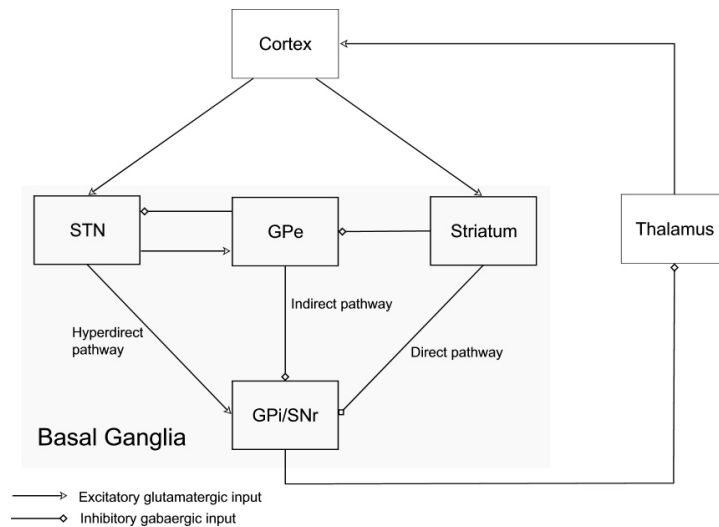


Figure 3: Schematic of corticostriatal pathways

Figure 3 Schematic of corticostriatal pathways illustrating a single corticostriatothalamic loop. The main pathway is the direct pathway. Competing intentions in the cortical module are selected amongst in the striatum. The selected intention is passed to the output nuclei, the Globus pallidus interna (GPe) and Substantia nigra pars reticulata (SNr) and thence back to the area of cortex it originated from. The hyperdirect pathway provides a fast excitatory pathway from the cortex via the subthalamic nucleus (STN) to the output of the basal ganglia

to focus the action selection. The indirect pathway provides a slower excitation to the output of the basal ganglia from the striatum via the Globus pallidus externa (GPe). This two step inhibitory pathway leads to excitation in the GPi and acts to terminate the selection of the intention.

The hyperdirect pathway provides a route for cortical excitation to be passed to the output nuclei of the BG.

Physiological operation

The cortical module proposes a number of contesting intentions. These are held in check by the tonic inhibitory output of the GPi/SNr acting via the thalamus. The striatum acts to decide amongst the competing intentions using information from past rewards obtained in similar environmental contexts (see section 4.7).

The three pathways provide mechanisms for intention selection, control of the force of the release of the intention and duration of release of the intention.

Derived computational functionality

We assume that a central function of corticostriatal circuits is action selection (or more accurately, intention selection). Specifically, the corticostriatal circuits in TOSCA will act as a winner-take-all network to mediate between mutually exclusive intentions. The main computation is performed at the level of the striatum where the intrinsic membrane properties of the principal neurons provide the capability to differentiate between the expected reward from each of the competing intentions.

Systems

As previously discussed, projections from posterior to anterior cortex can naturally encode associations between actions/intentions and features of the state that suggest that action. Multiple different, and potentially, conflicting intentions can be activated in parallel and it will often be necessary to select among conflicting actions. The neuroanatomy of corticostriatal circuits make them particularly well-suited to this function and interactions between cortex and basal ganglia will be crucial in doing so.

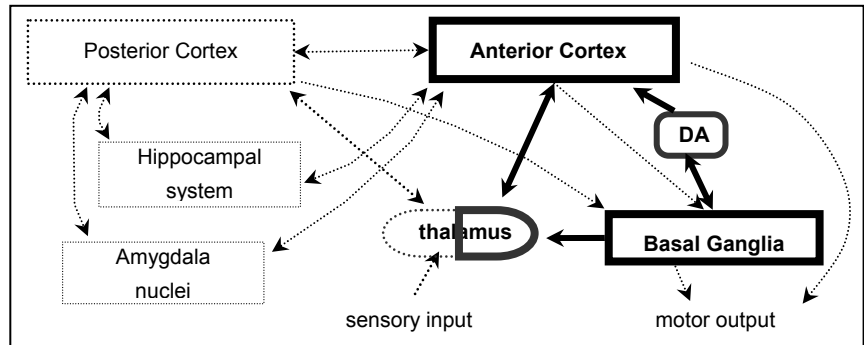
Interactions between this system and the dopamine system will also be crucial for learning in TOSCA. Specifically, when an action leads to unexpected reward, the value of that action in the current state/context will be increased (see section 4.7) by potentiating the cortical associations between the state features and the action representation. The corticostriatal action-selection system will be sensitive to these values, so that when that action is proposed in similar states in the future, its probability of being selected will be higher.

4.7 Dopamine reward circuits (reinforcement learning)

4.7.1 How the dopamine system modulates action contingencies

Anatomical structure

Dopamine producing neurons are located in two midbrain nuclei, the ventral tegmental area (VTA) and the substantia nigra pars compacta (SNc) (Figure 4). They receive excitatory input primarily from the pedunculopontine tegmental nucleus (PPTN) and prefrontal cortex and inhibitory input from the ventral striatum.



They project to the prefrontal cortex and striatum where they fire in a phasic fashion to release dopamine in response to rewarding situations (Romo & Schultz 1990, Schultz 1996).

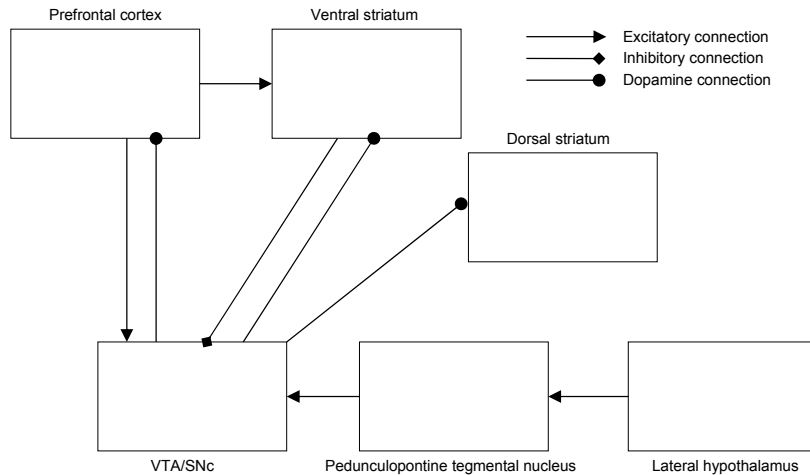


Figure 4: Schematic of corticostriatal pathways

Figure 4 Simplified schematic of the connections of the dopamine system. The nuclei containing the dopamine neurons, the ventral tegmental area (VTA) and the substantia nigra pars compacta (SNc) receive information about primary reward from the lateral hypothalamus via the pedunculopontine tegmental nucleus (PPTN). Neuromodulatory projections from the VTA/SNc are integral to learning from reward in the prefrontal cortex and striatum. The VTA also receive an excitatory input from prefrontal cortex and an inhibitory input from the ventral striatum.

Physiological operation

An unexpected (primary) reward elicits a phasic response in the dopaminergic neurons of the VTA/SNc. When a conditioned stimulus (CS) has been learned to reliably predict an upcoming reward, the time of response of the dopamine neurons shifts to coincide with the CS. These phasic releases of dopamine are utilized in the recipient structure to direct

learning so that intentions that are more likely to lead to reward are selected when the same environmental context is encountered in the future.

Derived computational functionality

Dopamine neurons have long been associated with reward learning and rewarded behavior, partly because of clear evidence of their key role in drugs of addiction (DiChiara, 1999), and because they are among the best targets for self-stimulation. The observation that the activity of dopamine cells in the monkey midbrain in reward-learning tasks closely follows the form of a key training signal in reinforcement learning (the temporal difference prediction error), is an important backdrop for TOSCA. In particular, temporal difference based RL methods will serve to modulate state-action associations by potentiating associations between clusters in posterior cortex (representing complex internal state information) and clusters in anterior cortex (representing internal and external action intentions).

4.7.2 Intrinsic Reward and its Neural Basis

Recent studies (Kakade & Dayan 2002, Dayan & Balleine 2002) have focused on the idea that dopamine not only plays a critical role in the extrinsic motivational control of behaviors aimed at harvesting explicit rewards, but also in the intrinsic motivational control of behaviors associated with novelty and exploration. For instance, salient, novel sensory stimuli inspire the same sort of phasic activity of dopamine cells as novel rewards (Schultz 1998, Horvitz et al. 1997}. However, this activation extinguishes more or less quickly as the stimuli become familiar. This may underlie the fact that novelty itself has rewarding characteristics (Montague et al. 1996).

The novelty-based release of dopamine onto one of its major targets, the striatum, causes both general psychomotor activation (Hooks & Kalivas 1994) and also specific exploratory or seeking behaviors such as approach that cause animals to engage with those novel stimuli. Approach of this sort is a Pavlovian response---it is like a pre-wired action inspired by novelty (and also reward prediction). Theoretical treatments (Kakade & Dayan 2001, Kakade & Dayan 2002) have directly related the dopamine activity with mechanisms for controlling exploration in the RL literature such as exploration and shaping bonuses (Sutton, 1993, Dayan & Sejnowski 1996, Ng et al. 1999) effectively completing the circle of interaction between computational, psychological and neural approaches. In TOSCA, we will explore a wider set of mechanisms by which animals control and benefit from exploration, using it to build sophisticated mechanisms for manipulating and exploiting novel environments. This wider set of mechanisms include the desire for mastery over one's environment and often leads to purposeful and sustained experimentation, as well as the motivation of an agent in a social setting to be liked by other agents (like-me) which leads to imitative behavior in social settings.

Various studies have also considered the neural basis of the assessment of novelty. Of particular relevance are two further neuromodulators, acetylcholine (ACh) and norepinephrine (NE), which are known to be involved in uncertainty and unexpectedness, and also to interact with the dopamine system. Theoretical treatments of these (Dayan & Yu 2003, Yu & Dayan 2002) focus on their roles in reporting specific sorts of uncertainty---

uncertainty arising from ignorance (which is what should drive exploration) and uncertainty arising from environmental stochasticity (which should not). The difference between these forms of uncertainty is relative to models of the environment, which form a key component of any theory of novelty. The ideas on ACh and NE are in their infancy; there is scope for a productive interaction between our explorations via TOSCA and future experiments and theory on the drives and effects of NE and ACh.

Derived computational functionality

The intrinsic motivations listed above will serve as mechanisms for providing internal reward to the agent and this in turn will help direct the agent's behavior during exploration and play both in the presence and absence of externally specified tasks. These internal rewards will lead to the learning of useful mental and physical skills in the form of *options* or abstract actions that in turn will become available to the reinforcement learning system in TOSCA as actions. This will allow an incremental buildup of a hierarchy of useful cognitive and physical skills by the agent that would not be possible in the absence of intrinsic motivations.

Systems

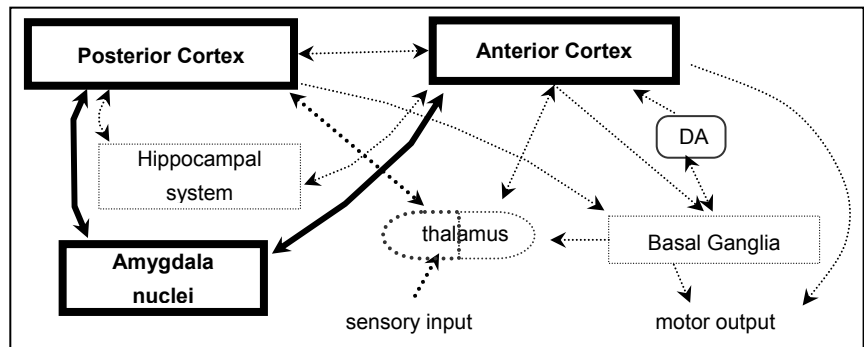
The dopamine system is tightly bound to the corticostriatal system, mediating learning in the prefrontal cortex and both divisions of the striatum. This system is also now known to provide neuromodulatory input to the hippocampal and thalamic systems.

4.8 Cortico-amygdala circuits (emotion, state-dependent storage & retrieval)

Anatomical structure

The amygdala formation is composed of multiple subparts typically grouped into the Medial group, the Central group, and the Basolateral group. The latter, forming the baso-lateral amygdala (BLA) is an

evolutionarily recent structure in contrast to the central and medial amygdala which are phylogenetically ancient. Central and medial amygdala nuclei are strongly connected to brainstem and hypothalamic structures and are implicated in visceral and hormonal modulation. See Figure 5 (left).






Figure 5: Amygdala nuclei (left) and interconnections with other structures (right)

The BLA is highly connected with portions of cortex including medial and lateral prefrontal cortex, sensory association cortex, as well as ventromedial frontal, rostral insular and rostral temporal cortical areas, and the medial thalamus and ventromedial basal ganglia. (See Figure 5 (right)). Connections from amygdala to cortex have recently been confirmed to preserve topographic organization (Amaral et al., 2003; Alheid 2003; Price 2003).

Physiological operation

Basolateral amygdala's topographic connectivity with anterior cortical regions is capable of "priming" or activating at a subthreshold level regions in anterior cingulate cortex, prefrontal cortex, and orbitofrontal cortex among others. Reciprocal activation is in evidence; i.e., amygdala and cortex activate each other (constrained by their topographic projection patterns).

Derived computational functionality

Analyses currently in progress suggest computational utility of cortico-amygdala circuitry; these include the reciprocal physiological priming effect described above, learning in both cortex and amygdala, and interaction between amygdala and hypothalamus. This work is currently being written up for publication (Hearn and Granger, in prep) and it is anticipated that the corresponding analyses will be included in the TOSCA architecture.

Systems

Cortico-amygdala circuits are integrated with cortico-striatal circuitry; these circuits have effects on behavior from sensory recognition to motor function, as well as on episodic memory storage and retrieval.

5. Key emergent functional properties: Representation & Control

The previous section lays out our vision for TOSCA at the level of brain systems and circuits. It explores the physiology we are trying to capture in TOSCA as well as the low-level computation being performed in individual brain systems and in brain circuits. However, it is down at a level where it is often difficult to see how human-level behavior emerges from these components and their connections.

Two primary features of the design of TOSCA are its representational system and its control system. The next two sections present initial overall views of how the underlying mechanisms

presented in section 4 operate together to generate an overall control mechanism to engage its environment, and to learn rich representations about that environment. Learning permeates the operation of the TOSCA system: the system is continually learning and cannot help but learn, and thereby builds up representations from combinations of perception and prior knowledge, as well as building up control knowledge.




Figure 6: Depiction of functional organization of TOSCA

Figure 6 depicts how the neural systems described throughout section 4 are functionally organized to produce useful behavior. The circled numbers correspond to subsections in subsection 4 where the underlying circuits that support the functionality are described. As can be seen in the bottom of the figure in the expanded diagrams, a given *functional* module in the figure (such as sequencing: 4.3, or intrinsic reward: 4.7) are realized by circuits that involve multiple brain systems.

The figure is organized on the left with the major contributors to the active internal state, which itself is organized vertically in the middle of the figure. Perception from various modalities (vision, auditory and touch are shown) provide low-level input [4.1]. Clustering [4.2] and sequencing [4.3] work from low-level perception, as well as other available internal state structures, to create more abstract structures that in turn become available for further clustering and sequencing as well as input for control (on the right side of the diagram) [4.6, 4.7]. The control system [4.6] consists of many parallel loops that work across both external actions (the motor system) and internal “mental” operations, including goal selection, persistence of working memory and goal structures, internal simulation, selective attention, and retrieval from episodic memory [4.5]. The reward system sends signals to the action selection modules (4.6) to tune performance through reinforcement learning.

5.1 Representation and memories: Multi-modal sequences of clusters and self-organizing maps

General properties of memories

The deceptively simple operators of sequences and categories, and the resulting data structures (hierarchically nested sequences of categories), interacting with special-purpose structures such as hippocampal, amygdala, and striatal formations give rise to the complete set of internal “knowledge representations” that occur in the TOSCA architecture. This surprising finding is in a way at the core of the TOSCA effort: it is a discovery of how advanced complex behaviors can be constructed from apparently simple interacting components.

Of particular importance is the emergence of interactions in what can be termed the primary architectural loop in TOSCA: the cortico-striatal loop. This set of circuits accounts for the vast majority of all the “real estate” in the entire system. Its behavior can be succinctly summarized thus: the representations generated by cortico-cortical systems can function as internal representations or models of states, which in turn can be tested and adaptively modified via reinforcement learning in cortico-striatal loop interactions. The resulting “adaptive exploitation” enables the construction of large and elaborate internal representations, and fits between those representations and the environment, via these basic powerful brain circuit mechanisms.

A note on learning via long-term potentiation (LTP)

In TOSCA, memories are stored via synaptic LTP, which operates via a set of well-worked out and extensively published and replicated rules and mechanisms. Many of these are unfamiliar to the field of psychology, and yield unusual memory effects in the architecture, all consistent with observed psychological phenomena.

Temporary memory:

Initial storage makes initial changes to synaptic weights. (Initial memory)

Consolidation:

If no new signals address those same synapses (storage sites) within the next 15-30 minutes (the synaptic consolidation period), those changes become permanent, i.e., irreversible. (Permanent memory).

Reversal:

If interfering signals do arrive at these storage sites within the synaptic consolidation period, the weight changes can be reversed. This can result in entirely forgetting or just “shaping” / altering the stored memory.

Elaboration:

Once stored, the memories can be internally accessed (via cortico-hippocampal loops) and give rise to internal practice and elaboration. In particular, new memories can be created to elaborate the initial memories, linking the initial memories to additional related memory items (e.g., seeing one car can become related to having driven in either that car or another car, etc), both enriching and altering the memories. This occurs via new synaptic recruitment and storage at new additional sites.

Thus all memories in the proposed architecture begin as temporary memories and can either become permanent (and possibly elaborated) or reversed (erased). Memories are stored where they are sensed or acted. There are no separate “locations” for memories of different durations.

Emergence of (multi-modal) categories = internal grammars

Two features of brain circuits past the sensory periphery are notable:

- i) circuits for different modalities (e.g., vision, audition) are remarkably similar (though not always identical; some of the gradient differences will be discussed separately); and
- ii) the majority of circuits receive inputs from multiple modalities. Thus communication among cortical regions consists of a single, shared, cross-modal internal representation language, regardless of the particular information being conveyed.

Individual cortical regions compute clusters (i.e., similarity-based categories) and sequences (chaining), via different components of their intrinsic circuitry. These two components, interact to produce sequences of categories (see Rodriguez et al. 2004). The output of one thalamocortical circuit is input to others with identical or near-identical structure; these thus produce sequences of categories of sequences of categories ..., effectively nesting the product of one “level” of processing into downstream processing products.

Successive nesting creates increasingly deep hierarchical “trees” of sequences of clusters. (Feedback from downstream to upstream regions participates actively in this process; partial activation of a downstream region has the consequence of increasing the probability of response of its potential upstream input constituents, acting in effect like “expectations” that those inputs will occur.)

These cortical mechanisms interact with hippocampal time dilation and contraction, amygdala “togglings” of salient features, and striatal reinforcement learning in cases of relevant feedback. Together the system produces incrementally constructed and selectively reinforced hierarchical representations consisting of nested sequences of categories (Granger 2006).

Figure 7 is an abstract illustration of successive stages of a representation so constructed. Initial simple input features (e.g., visual spots or edges; auditory frequencies or formants) transduced by front end mechanisms are learned by earliest, specialized stages (denoted in the figure by single letters A, B, etc). Their encoded outputs are input to downstream structures which learn clusters

(categories of similar inputs) and sequences of clusters; further downstream regions learn sequences of clusters of sequences of clusters, and so on.

Each downstream region, depending on its pattern of connectivity with its inputs, may exhibit a “bias”, preferring inputs with particular characteristics; these are genetically programmed and little is yet known of their layout, though work in quantitative neuro-anatomy is advancing knowledge in this realm. In TOSCA we will assume the existence of such biases, which cause different cortical regions to become increasingly “specialized” via learning for the particular feature combinations that they are most likely to successfully “compete” via lateral inhibition.

In practice, it would be prohibitively expensive computationally to learn all such combinations of features, but combinatorial explosion is avoided by two primary mechanisms:

- i) Bias: Of all the possible combinations of features that could occur, only some actually do, and, as just mentioned, some combinations are preferred over others;
- ii) Competition: With learning, oft-traversed regions become increasingly strengthened and, via lateral inhibition of neighboring regions, become what may be thought of as “specialists” in certain types of inputs, competing to respond.

Due to the described architectural arrangement, early upstream areas tend to respond to generic features and simple feature assemblies, but downstream regions respond with increasing selectivity to only specific assemblies, typically those that occur as patterns within oft-seen stimuli.

As a concomitant, further downstream regions should be expected to selectively respond to larger or longer patterns, both in visual and auditory domains. As most visual inputs consist simply of different arrangements of the same sets of primitive input features, it is expected that patterns of brain activation should be extremely similar in response to many different visual inputs, but that the similarity of those brain activation patterns ought to correspond to the similarity of their inputs, that is, activation patterns ought to be more similar for similar inputs, and more different for different inputs.

Moreover, if cortical regions are competing to respond to a given input, they should exhibit “category boundaries,” that is, the responses to images within a category (e.g., faces versus houses) should be more similar to each other than the images themselves are. Put differently, even highly different faces are likely to generate very similar cortical response patterns, whereas the similarity between any face and any house (as long as it is not a house that looks like a face!) should be more different than any two faces or any two houses.

These three sets of predictions from the model (distributed representations, similarity of patterns, and category boundaries) turn out to be controversial: depending on the analysis methods, neuroimaging studies have been used to support a number of still-conflicting hypotheses. TOSCA’s architectural design, as described, is concordant with some of the most prominent findings, in which distributed, overlapping patterns occur in response to images of, say, faces vs houses; more similar inputs tend to generate more similar responses; and responses to images within perceptual categories are more similar than responses to images across categories (Haxby et al. 2001; Pietrini et al. 2004; Furey et al. 2006; Hanson et al. 2004).

The representations thus far described bear some resemblance to long-term semantic memory: they are

- permanently stored,
- contain (learned) relations among components,
- tend to refer to categories and abstracts rather than individuals, and
- lend themselves to representations of generic types, e.g., “letters,” “speech sounds,” “jeeps.”

Characteristics of grammars (sequences of categories)

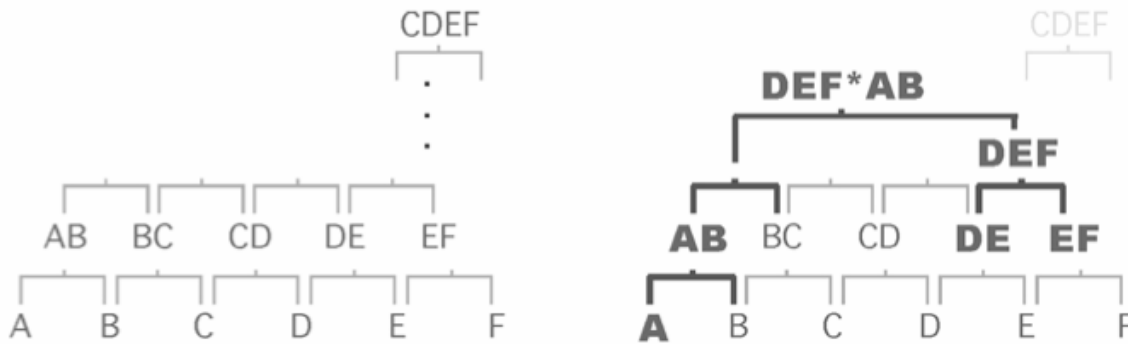


Figure 7: Illustration of hierarchies constructed by telencephalic architecture. Initial features generate successively nested sequences of categories of features (left). Additional exposure eventually (right) selectively strengthens sequences that recur (e.g., AB), weakens those that do not (e.g., CDEF), and constructs new sequences of categories as they occur and recur (e.g., DEF followed by a category that may include any of A-F (denoted here by a *) followed by AB).

The emergent data structure of the telencephalic system, statistically learned nested sequences of categories (as illustrated in Figure 7) is a superset of the structures that constitute formal grammars. The nested sequences of clusters are equivalent to ordered sequences of “proto-grammatical” elements such that each element represents either a category (in this case a cluster) or expands to another such element (nesting), just as grammatical rewrite rules establish new relations among grammatical elements.

Learning of the model’s representations as thus far defined (nested sequences of categories) constructs one type of semantic network referring to categories of objects, including relations among their internal parts (e.g., the hood, windows and trunk of a car).

Still to be specified are representations of a kind often occurring in such specifications – e.g., abstractions of relations (“in front of,” “above,” “containing”). These arise in a way compatible with hypotheses of “simulations” (Barsalou et al., 1999), i.e., learning specific instances in which objects are in the relation (a plate above a table, a hand above a paper, a window above a desk) generates not just representations of and among the particulars, but also abstract hierarchical representations of the relations themselves, which in turn become applicable to new inputs (a plane above a mountain) not previously seen.

The system extracts feature subsets as it learns, and generates regional cortical “specialists” as described earlier. Physical arrangements of objects in which one is higher in the visual field than

the other, for instance, lead to (relatively early) specialists that characterize the relationship between them. This relationship comes to have the verbal associations “above” and “below” (among others), and these relations come to internally define the corresponding abstract relations. This is also an area of still-active study in the architecture.

Emergence of high-level cognitive representations

The incremental nature of the “nested sequences of categories” data structure enables it to grow in function, simply by adding new copies of telencephalic thalamo-cortico-striatal-limbic loops – this functional growth corresponds to the incremental addition of “rules” acquired by the grammar. As more telencephalic “real estate” is added, the data structures that are constructed correspond to both longer and more abstract sequences, due to iterative nesting. Even regions of telencephalon with identical (or nearly identical) computational function nonetheless receive inputs from different sources, thus changing the feature combinations on which they operate (but see Galuske et al. 2000; Preuss 1995; 2000).

Proceeding “downstream” through the architecture, the outputs of one area are input to the next area. Successively more complex data structures should emerge, capturing increasingly complex representational concepts. Thus differential branching pathways through the architecture come to “specialize” in different functional realms.

Topics of ongoing study in the architecture concern the emergence of representational abstractions much-studied in psychology and in artificial intelligence, such as type-token distinctions (e.g., between “car” and “this car”), which enable distinguishing between individuals and categories. Initial study indicates that cortico-hippocampal interaction plays a role in this process, enabling the generation of different specifiers, qualifying cortico-cortical representations.

As mentioned, of the large set of all possible assemblies of features, only a small subset seem to be readily learned by biological organisms; there apparently exist species-specific biases that shape animals’ (including humans’) interpretations of various inputs. For instance, in response to very little data, humans will interpret certain coherent point-source motions as biological motion (e.g., when lights are affixed to the limbs of people moving in an otherwise dark environment); will interpret many distorted inputs as face-like; will interpret many sounds as speech-like, and so on. It is assumed that these biases may arise from developmental pre-selection (via mechanisms to be discussed elsewhere) of some cortico-cortical pathways that will selectively respond to particular types of feature assemblies.

It should be emphasized that all of these growing representational traits are hypothesized to arise directly from the hierarchical sequences of categories representations as manipulated by cortico-striatal loops. From low-level sensory beginnings, the abstractions grow to encompass the apparently full range of high-level cognitive concepts. This will be a central topic of study in the architecture.

Emergence of language

It has already been seen that the primary internal representation, hierarchically nested sequences of categories, is a form of grammar, i.e., shares the characteristics of formal grammar systems, though as described it has been used thus far for representations of sensory and motor sequences, not typically associated with grammars. In the TOSCA architecture, then, all internal representations are couched in the formalism of grammars (of this specific type).

If this representational hierarchy grows large enough (a function solely of the space of cortico-cortical structure allocated), the resulting abstractions become symbolic descriptors. At this point, the already grammar-based representation becomes the internal basis for linguistic representation. In other words, in the TOSCA architecture, grammars do not arise abruptly in service of linguistic abilities – rather, grammars are present throughout, and language arises as the representational hierarchy grows sufficiently large.

As described earlier, far-downstream areas are assumed to come to identify increasingly abstract symbolic descriptors (see Tables 2, 3) that are statistically repeated in relevant situations. These include definitions of words as well as the abstract relations that underlie the words' meanings.

Figure 8 illustrates structures occurring in response to simple sentences (“John hit Sam”) as input. Construction of sequences (e.g., S11, “John” followed by “hit”), and categories (e.g., C21, “hit” and “kissed,” items that can follow “John”) are combined in successive downstream regions (n+1, n+2, etc.) to create “proto-grammatical fragments” corresponding to internal representations of linguistic structure information.

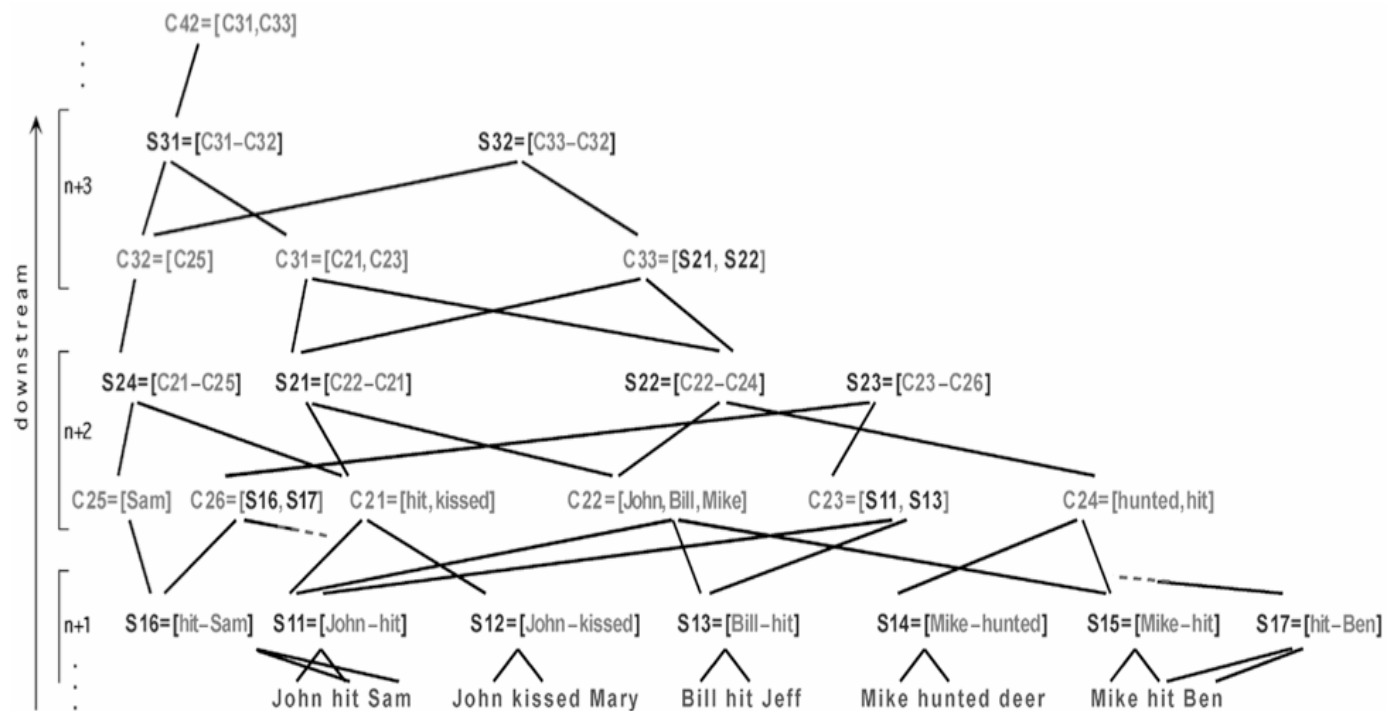


Figure 8: Nested sequences of clusters as sample proto-grammatical fragments educed from input strings.

It is worth noting that the generated structures can be used both a) in the processing of subsequent novel inputs and b) in the generation of arbitrary new strings that will conform to the rules inherent in the learned internal representational structures.

The resulting “generative” nature of the representations is worth emphasizing, addressing a crucial

aspect of linguistic grammars that can otherwise be absent from some purely input-processing or parsing mechanisms. A potentially infinite set of strings can be generated from the internal sequences of clusters, and the strings will be consistent with the internal grammar (see, e.g., Pinker 1999; Hauser et al. 2002; Fitch & Hauser 2004; Pinker & Jackendoff 2005).

It is also noteworthy that the grammar does not take the form typically adopted in attempts to formally characterize the syntactic structure of natural languages (such as English). The protogrammatical fragments capture regularities that are empirically seen to suffice for both parsing and generation, and have the structure to account for rule-like behaviors that characterize linguistic behavior. Research is currently in progress to study the formal relations between typical linguistic grammars, and protogrammatical fragments that are emergent from nested sequences of clusters.

(An additional characteristic of language that challenges researchers is the seeming effortless with which children learn language – readily contrasted even with the comparatively laborious training typically required for adults learning a second language. It is hypothesized that an innate bias related to sequences of categories of vocal utterances (speech) may lead (in larger-brained organisms) to a downstream bias for certain sequences of categories of assemblies of speech sounds (words). This may at least in part account for this much-studied but still elusive nature of innate language capacity; see Granger 2006).

5.2 Control: External motor resources and internal cognitive resources

In the previous section we described the emergent nature of representation in TOSCA—how both the present and the past come to be represented as multi-modal sequences of clusters at multiple levels of abstraction. In this section we describe how TOSCA exploits those representations to achieve adaptive, moment-by-moment control of both external motor effectors and internal cognitive states. The TOSCA theory of control is based on parallel loops of action selection contingent upon the representations described above and continuously modified by the intrinsic reward system. In the remainder of this section we first describe this general theory of control, followed by discussion of some of the specific control loops devoted to motor control and cognitive control.

5.2.1 General properties of control

TOSCA's general theory of control is based on three fundamental principles that have considerable biological and functional motivation: (1) fine-grained, parallel selection loops for both external and internal actions; (2) action selection potentially contingent upon multiple aspects of the internally represented state; (3) reinforcement learning of control realized by a rich intrinsic reward system. We now briefly summarize each of these in turn:

- *Fine-grained, parallel selection loops for both external and internal actions.* As we described above (Section 4.6) we assume that action selection is mediated by cortico-striatal loops, and more specifically, that multiple regions of frontal cortex represent competing intentions for action. There is growing neuroanatomical evidence that these loops are quite segregated [cite], so that the frontal-striatal system is best understood as consisting of many fine-grained selection loops operating in parallel. Functionally, this

organization is well-suited to support the real-time control of a motor system with many degrees of freedom (and much of the evidence for the segregation of frontal-striatal loops comes from detailed studies of mammalian motor systems), but as we outline in more detail below, it also naturally extends to the control of a cognitive system with multiple independent resources. Furthermore, in each specific case of cognitive control that we propose below, this extension is consistent with the existing biological evidence. Figure 6 above (right half) summarizes five major classes of control loops that we intend to model in TOSCA; each major control loop may be further broken down into separate finer-grained loops as described below.

- *Action selection contingent upon multiple aspects of the internal state.* One of the hallmarks of human cognition is its ability to adaptively exhibit arbitrary and novel behavioral contingencies. In prominent symbolic computational models of human cognition such as ACT-R and Soar, production rules play an important role in supporting this flexibility. More specifically, a critical property of production rules is that they allow action selection to be contingent upon any arbitrary features of the internally represented state (via the patterns in the "condition" side of the rules). Such functionality is directly supported in TOSCA by the massive inputs into frontal cortex from multiple posterior and anterior brain regions; this connectivity pattern and the neuro-anatomical evidence for it was described in Section 4.6. Figure 6 above (right half) depicts this broad contingency: each of the control loops starts with the activation of a set of *potential actions* that may be triggered by any aspect(s) of the internal state. These potential actions (or *intentions*) are represented in specific distinct frontal regions that participate in the segregated action loops described above.
- *Reinforcement learning of control.* Although the learning of behavioral and cognitive control in TOSCA ultimately depends on the interaction of multiple learning mechanisms in the architecture (including episodic encoding, clustering, and sequencing), the *direct* basis for learning control is reinforcement learning (RL) as realized by TOSCA's intrinsic reward system. We believe that much of the power of our proposed architecture will derive from the interaction of RL with the representational capacities of the system described above in Section 5.1. The functional neuroanatomy of this reward system was described in Section 4.7; the parallels to abstract properties of algorithms for reinforcement learning are well known (Shultz et al. 1997). The specific properties of RL in TOSCA are as follows:
- *Intrinsic reward.* All reward in the system is internally generated. "External reward" is translated from a sensation into a form of internal reward – there is no direct line from the environment to a reward signal. The internal rewards include intrinsic motivations or drives such as: *novelty*, *mastery*, and *exploration*. (See Section 4.7.2 above for the neural bases for these drives). These drives interact to determine the nature of both exploratory and task-driven behavior. The computational implications of this system for the nature of exploratory learning are significant and we draw them out in more detail in the next section below.
- *Exploitation of generalizations admitted by cortical representations.* In the previous section we described two critical ways that generalization is supported by cortical representations: via the emergence of *hierarchical clustering* realized by thalamocortical

loops, and via the self-organization of cortex into *maps* that support immediate similarity-based generalization. TOSCA's RL system will adaptively exploit these generalizations: contingencies (more specifically, *state-intention associations* as described above) based on useful abstract categories will come to be reinforced often and thereby strengthened. In this way, the system will learn to recognize abstract features of the state that are particularly helpful in determining which actions lead to intrinsic reward. Furthermore, because the representations of action intentions are themselves clustered and sequenced, the system may also acquire abstraction action plans that are also reinforced according to their intrinsic reward.

- *Exploitation of predictions admitted by thalamocortical sequence learning.* The previous section also described the sequence learning that continually operates over the learned categories; this sequence learning provides the functionality of *prediction* at multiple levels of abstraction (both semantic and temporal). This predictive capacity may then be exploited by the RL/reward system to significantly speed up learning in at least two different ways. First, the predictions as part of the internal state make the environment more observable to the agent. This can be a significant benefit because hidden state or equivalently partial observability significantly slows RL down (Singh et al. 2004). Second, the predictive capacity embodies an evolving probabilistic model of how the world evolves and can be used by the RL system to do "offline" learning or planning of state-action values which can lead to far better performance with far less "online" experience (Sutton 1990).
- *Parallel/cooperative reinforcement learning.* Each of TOSCA's segregated, multiple control loops is independently modulated by the intrinsic reward system. This structure makes TOSCA an instance of the more general class of *parallel reinforcement learning* systems, in which multiple, collaborating control systems interact through some shared state to maximize a shared reward signal. The key computational feature of such systems is that the explosive combinatorics of all the possible action combinations remains *implicit*: the control loops remain segregated, but cooperative behavior nevertheless emerges because each control loop adapts in the context of the behavioral consequences of the other co-adapting loops.

5.2.2 The role of intrinsic reward in shaping control

TOSCA's intrinsic reward system maps the rich internal state available to the agent to rewards that capture task-independent motivators (cf. Section 4.7) such as novelty, surprise, exploration, mastery (over environment), and like-me (by other agents in environment). These internal rewards lead the RL system in TOSCA to engage in exploration, play, and other behavior in the absence of explicit external reward (provided by some human specified task for example). As the agent engages in this behavior, the RL system learns a policy or rules of behaving that are captured in the form of the RL notion of options (Sutton et al. 1999). Options are temporally abstract actions that achieve subgoals and capture the intuitive notion of skills. Note that these skills could involve both external physical actions, for example an option could be about manipulating and mastering a physical object, as well as internal mental actions, for example an option could be about maintaining a particular episode of past experience in the internal state of the agent. These options or skills once learned become available as primitive actions to the agent

and can then be chosen by the action generation and selection mechanism in the same way as pre-wired actions can be chosen by the RL system. Thus, more complex skills can be learned that use skills learned earlier as components. This allows the agent to incrementally learn a hierarchical set of skills that lead to increasing competence over its environment (Singh et al. 2005) and this in turn makes the agent far more efficient at learning to solve externally specified tasks than would be possible without the internal reward based RL.

5.2.3 Control of motor system with multiple degrees of freedom

As discussed above there are multiple control loops in TOSCA and these loops allow both fine-grained control of individual degrees of freedom as well as coordinated control over multiple degrees of freedom in the form of motor routines. Some basic motor routines or skills will be pre-wired into the agent but many will be learned using the intrinsic reward based RL system outlined above. For example, driven by the internal motivation to achieve mastery over an object in its environment the parallel RL system described above would learn a complex motor routine or skill that orchestrates multiple control loops over time to reach for and manipulate that physical object as well as maintain internal state needed to accomplish the manipulation. Once learned, these skills that coordinate multiple control loops become available as primitive action choices to the parallel RL system leading to even more complex and richer hierarchical control of the motor system.

5.2.4 Control of cognitive resources

The rich multi-modal representational and memory systems described in Section 5.1 provide more than the basis for overt behavioral contingencies: they are themselves cognitive resources under adaptive control. The nature of this cognitive control ranges from the modulation of representations of current perceptions to the use of multi-modal imagery to simulate novel dynamic situations. In short, the system has control over aspects of its own internal state. This is a critical computational feature because it allows the system to move beyond reactivity to the kind of open-ended behavior that depends on arbitrary aspects of the past as well as the ability to flexibly project into the future. Figure 6 above depicts five critical classes of cognitive control loops which we briefly describe below. In all cases, what mediates top-down control is associations from frontal cortical areas (both to posterior regions and other frontal regions) that represent the selected cognitive actions. These frontal action representations are in turn contingent upon internal state and the frontal-striatal selection mechanisms described earlier.

Control of attention. The term "attention" has many meanings in psychology and cognitive neuroscience; we use it here to refer specifically to mechanisms of *selective activation and enhancement of perceptual representations*. For example, if it is adaptive for the system to attend to color features in certain situations, this may be accomplished by the selective enhancement of color features in lower-level perceptual representations which then bias the resulting higher-level categories to be more sensitive to color distinctions. Such changes in categorization then affect what new actions (internal and external) are selected next. In this way even relatively low-level attentional modulation has qualitative effects on behavior; the reward system reinforces those attentional contingencies that have positive effects.

Control of multi-modal working memory. Most tasks require the integration of information

(either perceptual inputs or intermediate computational products) over time. There is considerable neural evidence that the persistence of such information over relatively short time periods (seconds to tens of seconds) depends critically on cortical representations (independent from the hippocampal subsystem). In TOSCA this persistent information is represented in the same higher-level perceptual posterior areas where the stimuli were originally processed. The information is maintained by excitatory connections from persistent (attractor-based) representations in prefrontal cortex [cite]; these prefrontal regions represent the "action" of maintaining a specific type of perceptual input. Those maintenance action contingencies that lead to intrinsic reward are reinforced. In this way TOSCA will learn to maintain task-relevant information in the face of potentially interfering irrelevant stimuli [cite o'reilly].

Control of multi-modal imagery and "simulators".

Another important emergent feature of TOSCA's architecture is the ability to reactive modality-specific systems under cognitive control. Substantial evidence now supports this hypothesis (Barsalou (2003b); Barsalou, Simmons, Barbey, and Wilson (2003); Barsalou, Niedenthal, Barbey, and Ruppert (2003); Niedenthal, Barsalou, Winkielman, Krauth-Gruber, & Ric (2005); Martin (2001); Pulvermüller (1999); Thompson-Schill (2003); Smith and Gasser (2005); and Thelen (2000).) The frequent use of modality-specific systems by higher cognitive processes provides interesting new functionality. Barsalou (1999, 2003a, 2005) explains these emergent capabilities in terms of *simulators*, which can be naturally implemented via the sequences-of-clusters mechanism that underlies the TOSCA architecture.

Much research has shown that categories have statistically correlated features (e.g., *wheels*, *steering wheel*, and *engine* for *cars*; McRae, de Sa, & Siedenberg, 1997). Thus, encountering different instances of the same category should activate similar neural patterns in feature systems (cf., Farah & McClelland, 1991; Cree & McRae, 2003). Furthermore, similar populations of conjunctive neurons in the brain's association areas—tuned to these particular conjunctions of features—should tend to capture these similar patterns (Damasio, 1989; Simmons & Barsalou, 2003). Across experiences of a category's instances, this population of conjunctive neurons integrates the modality-specific features of a category, establishing a distributed multi-modal representation of it. Thus, a simulator is a distributed circuit that includes brain areas for representing modality-specific content and integrating it.

Conceptually, a simulator functions as a type: It integrates the multimodal content of a category across instances, and provides the ability to interpret later individuals as tokens of the type. Consider the simulator for the category of *cars*. Across learning, visual information about how cars look becomes integrated in the simulator, along with auditory information about how they sound, somatosensory information about how they feel, motor programs for interacting with them, emotional responses to experiencing them, etc. The result is a distributed system throughout the brain's feature and association areas that accumulates modal representations of the category.

Control of declarative/episodic memory retrieval. Although episodic encoding may be automatic, there is substantial psychological and cognitive neuroscience evidence that episodic memory retrievals are often under deliberate control. In TOSCA this controlled retrieval is realized via a combination of several mechanisms. Retrieval starts with the assembly of *retrieval cues* accomplished by the control of multi-modal working memory and imagery (see above; mediated primarily by frontal-posterior projections). These representations serve as cues by

activating conjunctive episodic representations in the hippocampus via connections from posterior cortex to the hippocampus (see Section 4.5). But the hippocampus must be biased to process its inputs as cues for retrieval rather than new episodes to encode; this biasing happens via projections to the hippocampus from frontal regions that represent the specific intended cognitive action of retrieval [for precedents in the literature for such controlled retrieval see O'reilly, Eichenbaum]. Again, those contingencies for retrieval actions that lead to intrinsic reward are reinforced, so that the system learns how to make effective use of its own episodic memory system.

Control of working memory for goals (abstract intentions). In the moment-to-moment control of behavior, the system faces the immediate problem of selecting the next best action to perform (across its multiple control loops). But as described above, this action selection may be contingent upon any feature of the internal state, including abstract features. Given that internal state is also under control—and crucially state that may *persist* for seconds to tens of seconds (see above)—the system can manipulate persistent internal state in order to direct its own behavior. In short, TOSCA will have the capacity to establish specific *goal representations* in working memory, and to learn when it is advantageous to do so. Functionally, such representations differ from the multi-modal working memory representations described above because they are exclusively about the *control* of future behavior and not about the temporary maintenance of relevant perceptual or imaginal information. Furthermore, the clustering over action sequences provides a suitable representational vocabulary for hierarchical goals because they are abstract intentions that already have associations with their specific constituent sequenced actions.

5.2.5 Social learning

In humans, learning rarely occurs save in the presence of other humans; from early development through adult stages, humans act predominantly in settings where other humans interact, largely via language. Interaction conditions include both direct learning from being told, as well as mixes of partially-guided and partially-independent discovery. In each case, balances must be struck between learning on one's own (via internal rewards and reinforcement learning) and learning from the social environment. Following are some of the ways social interaction will help speed up the reinforcement learning of skills and tasks. 1) The visual attention of the agent could be directed by the human guide towards socially salient or task/goal relevant stimuli, 2) The human guide should be able to demonstrate or suggest actions (using language) that coupled with the intrinsic reward for being liked would lead to the agent to explore those actions, 3) The human guide may provide explicit reward to the reinforcement learning system inside TOSCA, and 4) The human guide may provide subgoals by pointing out landmark states or other novel states relevant to the overall task faced by the agent. All of these forms of social interaction directly impact the reinforcement learning system by either providing internal or external reward or suggesting salient states to pay attention to or providing salient actions to imitate. Together, these can significantly reduce the amount of data needed by the agent to learn competence over its environment.

5.2.6 Language and control

The previous section highlighted the important role of language in social learning. Here we

consider one specific avenue through which language has a dramatic impact on cognition and behavior: instruction taking. The ability to take instructions is computationally significant because it means that the system is *immediately taskable*, providing a way to move past the incremental adaptation that is characteristic of reinforcement learning. In this section we provide a sketch of how the architectural mechanisms in TOSCA will support instruction taking, focusing on the situation where the system acts immediately upon its received linguistic input. (Instruction taking situations involving greater temporal separation of the instruction and behavior would tap into the same mechanisms described here, but would additionally involve the long term declarative memory system).

At a high level, instruction taking can be decomposed functionally into (a) *comprehension* of the linguistic input, which yields representations that are (b) *interpreted* to produce behavior. (These separate functions need not be strictly staged but may be tightly interleaved in the process of incremental understanding and behaving). The representational and functional capacities of TOSCA described above are sufficient to accomplish these functions and, crucially, to accomplish these functions with novel linguistic inputs. Briefly, the process would work as follows:

- *Comprehension* happens as the linguistic input is incrementally processed and given a hierarchical representation as described above in Section 5.1 (which outlined how certain aspects of linguistic grammars may emerge from clustering and sequencing). The critical feature is that the linguistic input is represented at multiple levels of abstraction in a way governed by the (learned) grammar of the language, making explicit the structural (and thus indirectly the semantic) relations among the constituents parts of the input.
- Task-critical aspects of these representations are *selectively maintained* as an assembly of cortical representations in the frontal inferior and posterior language areas so that they may guide future behavior (see Section 5.2.4 above on adaptive control of working memory). This provides an important link between comprehension and interpretation.
- The *interpretation* of the instructions happens as these linguistic representations function as goals upon which action selection (in the multiple parallel loops) is contingent. Such linguistically-driven action selection is learned via the same reinforcement learning mechanisms described earlier. A crucial aspect of this learning is that it exploits the abstract generalizations admitted by the hierarchical representations (see section 5.2.1). This generalization, coupled with the fact that the interpretation takes place incrementally, will give rise to an interpretive skill that will transfer immediately to novel linguistic inputs that share critical structure. For example, the interpretation of simple instructions taking the form *perform a specific action upon an object* will consist of separate acts of orienting and grasping etc. that will be useful for a wide range of different specific actions. The most-often rewarded contingencies for such initial orienting will thus tend to abstract away (via the clustered representations) from the specific instructed actions. Similarly, contingencies for aspects of the control to accomplish the instructed action may abstract away from the specific linguistic label used to identify the object, because such abstract contingencies are reinforced more often.

Thus TOSCA's ability to take instructions will arise from an interaction of the emergent, generative grammatical representations described above and the reinforcement-learning-based interpretation of linguistic representations that exploits the learned abstract grammatical categories. The reinforcement learning of the interpretive skill is part of the gradual, incremental

process of language acquisition, but it crucially yields immediate taskability. In short, language comprehension and use is a procedural skill operating on special types of representations, but acquired via general processing principles. This novel approach to instruction taking has the virtue that learning permeates every aspect of the processing, and we believe it will provide a compelling demonstration of the cognitive power that emerges from the interaction of TOSCA's basic architectural mechanisms.

6. Conclusion

The goal of this paper is to describe our design for a new cognitive architecture based on the brain: TOSCA. With many cognitive architectures, a design is a straight forward description of component modules (such as procedural memory, working memory, bottom-up sensor processing), and a straightforward mapping of components onto high-level functionality. But in designing the brain, evolution took an interesting turn. Instead of having a parallel decomposition of functionality and structure, it built on a more primitive set of computational components that we presented in Figure 1 of Section 3. These primitives are tightly interconnected and form circuits (Section 4), so that functionality emerges from the interactions of these multiple circuits (Section 5).

This organization starts to answer the question as to where is the “magic in human cognition” – it is not in any one module that AI hasn’t yet discovered, but it is in choosing the right set of building blocks and the connections between them. But above and beyond the static organization is the dynamics of the system where learning is ubiquitous. For learning to be successful, the right information must be made available, and this organization brings the information to the right places. For some components the learning is mostly bottom-up, as the system learns statistical regularities of its environment (as in learning clusters and sequences). For others it is associational where learning brings together co-occurring sensations across modalities. And for action, intrinsic reward drives the learning of control across not just external actions, but the system’s control of itself. This is the path for TOSCA, and the path to a new generation of brain-based cognitive systems.

References

[In this draft, our list of references is incomplete and not standardized or sorted.]

- Alexander GE, DeLong MR, Strick PL. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience* **9**, 357–38.
- Arbuthnott GW, MacLeod N, Rutherford A. The rat cortico-striatal pathway in vitro. *Journal of Physiology London* **367** (1985), 102P.
- Barsalou, L.W. (1987). The instability of graded structure: Implications for the nature of concepts. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 101-140). Cambridge: Cambridge University Press.
- Barsalou, L.W. (1989). Intraconcept similarity and its implications for interconcept similarity. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 76-121). Cambridge: Cambridge University Press.
- Barsalou, L.W. (1993). Flexibility, structure, and linguistic vagary in concepts: Manifestations of a compositional system of perceptual symbols. In A.C. Collins, S.E. Gathercole, & M.A. Conway (Eds.), *Theories of memory* (pp. 29-101). London: Lawrence Erlbaum Associates.
- Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, **22**, 577-660.
- Barsalou, L.W. (2003a). Abstraction in perceptual symbol systems. *Philosophical Transactions of the Royal Society of London: Biological Sciences*, **358**, 1177-1187.
- Barsalou, L.W. (2003b). Situated simulation in the human conceptual system. *Language and Cognitive Processes*, **18**, 513-562.
- Barsalou, L.W. (2005). Abstraction as dynamic interpretation in perceptual symbol systems. In L. Gershkoff-Stowe & D. Rakison (Eds.), *Building object categories* (389-431). Carnegie Symposium Series. Mahwah, NJ: Erlbaum.
- Barsalou, L.W., Niedenthal, P.M., Barbey, A., & Ruppert, J. (2003). Social embodiment. In B. Ross (Ed.), *The Psychology of Learning and Motivation*, Vol. 43 (pp. 43-92). San Diego: Academic Press.
- Barsalou, L.W., Simmons, W.K., Barbey, A.K., & Wilson, C.D. (2003). *Grounding conceptual knowledge in modality-specific systems. Trends in Cognitive Sciences*, **7**, 84-91.
- Cree, G. S., & McRae, K. (2003). Analyzing the factors underlying the structure and computation of the meaning of chipmunk, cherry, chisel, cheese, and cello (and many other such concrete nouns). *Journal of Experimental Psychology: General*, **132**, 163-201.
- Damasio, A.R. (1989). Time-locked multiregional retroactivation: A systems-level proposal for the neural substrates of recall and recognition. *Cognition*, **33**, 25-62.
- Farah, M. J., & McClelland, J. L. (1991). A computational model of semantic memory impairment: Modality specificity and emergent category specificity. *Journal of Experimental Psychology: General*, **120**, 339-357.
- Kemp JM, Powell TPS. The corticostriate projection in the monkey. *Brain* **93** (1970), 525-546.
- Martin, A. (2001). Functional neuroimaging of semantic memory. In R. Cabeza & A. Kingstone (Eds.), *Handbook of functional neuroimaging of cognition* (pp. 153-186). Cambridge, MA: MIT Press.
- McRae, K., de Sa, V.R., & Seidenberg, M.S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, **126**, 99-130.
- Montague PR, Dayan P, Sejnowski TJ. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, **16** (1996), 1936–1947.
- Niedenthal, P.M., Barsalou, L.W., Winkielman, P., Krauth-Gruber, S., & Ric, F. (2005). Embodiment in attitudes, social perception, and emotion. *Personality and Social Psychology Review*, **9**, 184-211.
- Pulvermüller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, **22**, 253-336.
- Romo R, Schultz W. Dopamine neurons of the monkey midbrain: contingencies of responses to

- active touch during self-initiated arm movements. *The Journal of Neurophysiology* **63** (1990), 592-606.
- Schultz W. Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *The Journal of Neurophysiology* **56** (1986), 1439-1461.
- Schyns, P.G., Goldstone, R.L., & Thibaut, J.P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1-54.
- Simmons, W.K., & Barsalou, L.W. (2003). The similarity-in-topography principle: Reconciling theories of conceptual deficits. *Cognitive Neuropsychology*, *20*, 451-486.
- Smith, L., & Gasser, M. (2005). The development of embodied cognition: Six lessons from babies. *Artificial Life*, *11*, 13-29.
- Sutton, R.S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. Proceedings of the Seventh International Conference on Machine Learning, pp. 216-224, Morgan Kaufmann.
- Sutton, R.S., Precup, D., Singh, S. (1999). Between MDPs and semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence* *112*:181-211
- Thelen, E. (2000). Grounded in the world: Developmental origins of the embodied mind. *Infancy*, *1*, 3-30.
- Thompson-Schill, S.L. (2003). Neuroimaging studies of semantic memory: inferring “how” from “where”. *Neuropsychologia*, *41*, 280-292.
- D. E. Berlyne. A theory of human curiosity. *British Journal of Psychology*, *45*:180–191, 1954.
- D. E. Berlyne. Conflict, Arousal, and Curiosity. McGraw-Hill, N.Y., 1960.
- D. E. Berlyne. Aesthetics and Psychobiology. Appleton-Century-Crofts, N.Y., 1971.
- D. O. Hebb. The Organization of Behavior. Wiley, N.Y., 1949.
- G. Di Chiara. Drug addiction as dopamine-dependent associative learning disorder. *European Journal of Pharmacology*, *375*(1-3):13–30, 1999.
- G. Barto. Adaptive critics and the basal ganglia. In J. C. Houk, J. L. Davis, and D. G. Beiser, editors, *Models of Information Processing in the Basal Ganglia*, pages 215–232. MIT Press, Cambridge, MA, 1995.
- K. J. Friston, G. Tononi, G. N. Reeke, O. Sporns, and G. M. Edelman. Value-dependent selection in the brain: Simulation in a synthetic neural model. *Neuroscience*, *59*:229–243, 1994.
- S. Kakade and P. Dayan. Dopamine bonuses. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 131–137. MIT Press, 2001.
- S. Kakade and P. Dayan. Dopamine: Generalization and bonuses. *Neural Networks*, *15*:549–559, 2002.
- Ng, D. Harada, and S. Russell. Policy invariance under reward transformations: Theory and application to reward shaping. In Proceedings of the Sixteenth International Conference on Machine Learning. Morgan Kaufmann, 1999.
- P. R. Montague, P. Dayan, and T. J. Sejnowski. A framework for mesencephalic dopamine systems based on predictive hebbian learning. *Journal of Neuroscience*, *16*:1936–1947, 1996.
- P. Dayan and B. W. Balleine. Reward, motivation and reinforcement learning. *Neuron*, *36*:285–298, 2002.
- P. Dayan and T. J. Sejnowski. Exploration bonuses and dual control. *Machine Learning*, *25*:5–22, 1996.
- P. Dayan and A. J. Yu. Uncertainty and learning. *IETE Journal of Research*, *49*:171–182, 2003.
- S. Kakade and P. Dayan. Dopamine bonuses. In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 131–137. MIT Press, 2001.

- S. Kakade and P. Dayan. Dopamine: Generalization and bonuses. *Neural Networks*, 15:549–559, 2002.
- M. S. Hooks and P. W. Kalivas. Involvement of dopamine and excitatory amino acid transmission in novelty- induced motor activity. *Journal of Pharmacology & Experimental Therapeutics*, 269:976–988, 1994.
- J. C. Horvitz, T. Stewart, and B. Jacobs. Burst activity of ventral tegmental dopamine neurons is elicited by sensory stimuli in the awake cat. *Brain Research*, 759:251–258, 1997.
- W. Schultz. Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80:1–27, 1998
- W. Schultz, P. Dayan, and P. R. Montague. A neural substrate of prediction and reward. *Science*, 275:1593–1598, March 1997.
- R. S. Sutton, D. Precup, and S. Singh. Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112:181–211, 1999.
- J. Yu and P. Dayan. Acetylcholine in cortical inference. *Neural Networks*, 15:719–730, 2002.
- S. Singh, M. R. James, and M. R. Rudary. Predictive State Representations: A New Theory for Modeling Dynamical Systems In Uncertainty in Artificial Intelligence: Proceedings of the Twentieth Conference (UAI), pages 512-519, 2004.
- S. Singh, A. G. Barto, and N. Chentanez. Intrinsically Motivated Reinforcement Learning. In *Proceedings of Advances in Neural Information Processing Systems 17*, 2005.

Appendix: TOSCA Framework

The TOSCA Framework is a toolkit that will be used to build one or more specific TOSCA architectures. The Framework should be general purpose and flexible so it supports the design and experimentation on particular TOSCA architectures. This means the focus for the Framework is on generality and flexibility of infrastructure, while a particular architecture might be much more tightly constrained and make specific behavioral predictions.

To ensure that any TOSCA architecture makes strong biologically-inspired commitments, the TOSCA Framework imposes a set of constraints on the architectural design. First, all processing must be mapped to a specific brain region. This ensures that an executing TOSCA system always makes a formal commitment to how all simulated processing would map to biological processing. Second, all communication between modules, within the simulation, must conform to known biological constraints based on the current literature. If the connections made within a TOSCA simulation violate the known connectivity properties of regions within the brain, these constraint violations will be explicitly detected by the Framework.

In order to support the expression of these biological constraints, the processing within the Framework must have clearly defined boundaries. The proposal is to divide the elements of the Framework into two categories:

- Function Modules
- State Variables

Function modules perform processing and state variables are used for communication between modules. Modules receive inputs from a set of state variables and generate outputs to one or more state variables.

The critical constraints are that values can only be persistently stored within state variables and all state variables must be mapped to a brain region. Together, these constraints imply that for any model implemented within the TOSCA Framework we can determine how the model is connected and what commitments the model makes for regions of the brain where processing occurs and how those brain regions are connected. Those model commitments can then be reviewed against current knowledge of biological processing to determine the quality of the fit.

The relationships between function modules and state variables and how they map to biological processing are represented through three logical graph structures within the TOSCA Framework:

1. Functional Connectivity Graph
2. Brain Mapping Graph
3. Brain Connectivity Graph

Functional Connectivity Graph

This graph establishes a mapping of state variables to input and outputs of function modules. It defines the flow of data through the simulation.

For example, Figure A1 shows module M1 receiving inputs from state variables A, B and C and generating outputs for A and D. Module M2 receives input from D and generates output for C.




Figure A1: Functional Connectivity Graph

This graph is implicitly defined by the inputs and outputs from function modules. It is not explicitly represented in the Framework as a separate data structure, although we will have tools that can extract this graph for display or analysis based on examining the connections between modules and state variables.

State variables can also be used to provide private storage where only a single function module ever accesses the state variable. Function modules are potentially executing asynchronously and at a range of time scales.

Some key properties of this graph are:

- **Flexible replacement of modules and state variables**

A module or state variable can be replaced with a different implementation with minimal impact on other modules and variables. This supports experimentation by swapping in different implementations for parts of the overall architecture. It also supports incremental and distributed development of the entire architecture as placeholders can be used until a full component has been developed and can replace the placeholder.

- **State variables are used for all persistent data**

All data that persists from one cycle of the simulation to another (i.e. that is not temporary storage used in a calculation) should be represented by a state variable.

Some examples of state variables are:

- A frame buffer of video input
- Connection strengths in a neural network
- Variables used to control search
- Matrix or vector of floating point values
- Production rules (condition -> action rules)

This property relates to the brain mapping graph (below). Only state variables would be mapped onto brain regions. There would be no need to map function modules to brain regions explicitly as that mapping would be implicit in the state variables.


Brain Mapping Graph

The brain mapping graph establishes how function modules and state variables are mapped to different regions of the brain. In the current design only state variables are explicitly mapped to brain regions. Function modules are mapped to regions based on the state variables they use (Figure A2).

Unlike the functional connectivity graph, the brain mapping graph will be explicit. There will be a specific data structure in the Framework that defines how each state variable is mapped to a region of the brain.

The mapping:

- **Will be complete**
All state variables must map to some brain region. This ensures that there's an explicit mapping from a running simulation to the brain.
- **May include unspecified regions**
A state variable may have an undetermined location but this must be explicitly indicated in the mapping.



F

Figure A2: Brain Mapping Graph

Brain Connectivity Graph

The functional connectivity graph together with the brain mapping graph imply a brain connectivity graph (e.g. Figure A3).

That is, by examining how the state variables and function modules are connected together and how they are mapped to brain regions makes a prediction about how the brain regions are connected. This graph can then be tested against known constraints for how brain regions are actually connected. These constraints will be explicitly represented within the framework and the brain connectivity graph will be deduced automatically from the brain mapping graph and the functional connectivity graph.




Figure A3: Brain Connectivity Graph

Function Modules and Time

Function modules may execute asynchronously and at different time scales. Two communicating modules may be executing on a single machine or across a network between a cluster of machines. The TOSCA Framework provides an abstraction layer over these details, allowing each module (and set of state variables) to be defined in terms of its own temporal constraints, with the Framework handling the details of communication and synchronization between modules and variables.

Figure A4 shows an example of the capabilities provided by the Framework. Modules M1 and M2 are executing on a 1ms and 5ms time scale respectively. Every 50ms M2 posts a result to the state variable D which is consumed by M3. This synchronization is achieved through two different clocks – T1 is the master clock for the simulation and T2 is a local clock used only to synchronize the behavior of M1 and M2.

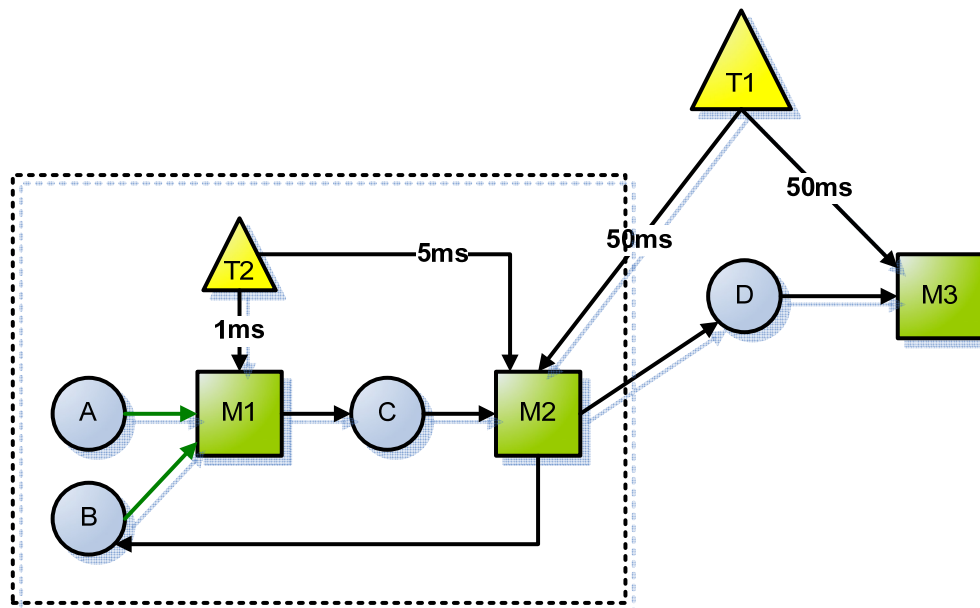


Figure A4: Local Clock Example

Each module expresses the constraints it is placing on the overall flow of time through the simulation by making calls an appropriate clock. Based on these constraints, the clock(s) determine when simulation time can advance and by how much before further processing must occur in a module. Each module is unaware of the complexity within the entire system and deals only in the local constraints that concern it (e.g. M2 reads input from C every 5ms and outputs to D every 50ms). Additionally, everything within the dotted line shown in *Figure A4* could be collected into a larger function module. This allows hierarchies of modules to be built and used in experimentation. For instance, two different implementations of a particular function module (the first a simple module and the second a complex collection of submodules, variables and clocks) could be tested against each other without modifying the rest of the simulation. The sample code below shows the detail of how a function module would appear within the Framework if written in C++ (the actual Framework will support multiple implementation languages):

```

while (!m_Quit)
{
    // Wait for a signal from the clock (if time driven)
    // or from an input changing (if event driven)
    WaitForNextSignal() ;

    // Process any events we were sent while sleeping
    ProcessEvents() ;

    // Get the current clock time. This value can't change until
    // we signal that it's ok for the clock to advance.
    Time time = GetClock()->GetTime() ;

    // Get the value of the inputs at the current time
    Value a = GetInput("A")->GetValue(time) ;
    Value b = GetInput("B")->GetValue(time) ;
    Value c = GetInput("C")->GetValue(time) ;

    // We've read our inputs so the clock can continue
    // but it can't go beyond the time when we generate output.
    // Changing the boundary time allows the clock to advance.
    GetClock()->ClockCannotAdvanceBeyond(this, time+4) ;

    // Perform a calculation based on the inputs (could be long time)
    Calc(&a, &b, &c, &result) ;

    // Generate output 5ms after this function module was triggered
    GetOutput("A")->SetValue(&result, time+5) ;

    // Set up our next triggering event
    GetClock()->RegisterWakeup(this, time+10) ;

    // Clock can now advance freely up to our next trigger event
    GetClock()->ClockCannotAdvanceBeyond(this, time+10) ;
}

```

The module initially waits for a signal (either a specific amount of time passing or an input variable changing) before beginning processing. It receives any notifications of new system events before reading the current input values for state variables A, B and C. While reading its inputs the simulation clock is kept at a known time, ensuring that the inputs remain valid. Once the module has read all of its inputs it signals that the clock can now advance but not beyond the point when the module will generate output. The module computes the output and posts the new value (in this case in a feedback loop to variable A, 5ms after it awoke). It then requests its next wakeup signal—in this example by asking the clock to signal it when 10ms have passed and releases the clock to move up to that new time.

This example shows a module that is *time-driven*, where it's processing is triggered by the flow of time. A very similar logic is used when the module is *event-driven*, where it's processing is triggered by the change of an input. The module registers in advance with the input variable to request notification when the variable changes state and then engages in similar processing to the time-driven example shown here.

The Framework maintains the actual location of all state variables, which could be on different physical machines, as well as determining the correct flow of simulated time within the constraints each module imposes. The specific location of a state variable can be changed without affecting existing modules and the details of how a module is implemented can be changed without affecting other modules as long as

its inputs and outputs are not significantly changed.

Implementation Flexibility

The Framework will allow modules to be written in a variety of different languages and then combined together to form a complete model. This is achieved by implementing the core functionality in C/C++ and then machine generating interfaces in other languages. We have other projects that have adopted this method and support 5+ languages in this manner. This cross-language capability allows each researcher or team to adopt the language that best supports their work and yet still create an integrated system.

The Framework will also support a range of runtime configurations, from a single process executing the entire simulation, to a series of processes using shared memory within a single machine and up to a cluster of machines each executing a part of the entire simulation. This flexibility is achieved by basing all of the communication on message passing between state variables and then abstracting over the details of how those messages are passed—e.g. by accessing a local pointer or by sending a message over a socket to a different machine. This abstraction allows the hardware infrastructure to scale up as the requirements of a particular task increase, without having to redesign modules or have different implementations for different runtime configurations.